Metrics, Connections, & Correspondence: The Setting for Groupwise Shape Analysis

Carole Twining[1], Stephen Marsland[2], and Chris Taylor[1]

[1] Imaging Science (ISBE), University of Manchester, Manchester, UK.
[2] School of Engineering and Advanced Technology (SEAT), Massey University, Palmerston North, New Zealand.

Abstract. This paper considers the general problem of the analysis of groups of shapes, and the issue of correspondence in that context. Many papers have been published on the topic of pairwise shape distances and pairwise shape similarity measures. However, most of these approaches make the implicit assumption that the methods developed for pairs of shapes are then sufficient when it comes to the problem of analyzing groups of shapes. In this paper, we consider the general case of pairwise and groupwise shape analysis within an infinite-dimensional Riemannian framework. We show how the issue of groupwise or pairwise shape correspondence is inextricably linked to the issue of the metric. We discuss how data-driven approaches can be used to find the optimum correspondence, and demonstrate how different choices of objective function lead to different groupwise correspondence, and why this matters in terms of groupwise modelling of shape.

1 Introduction

It is generally agreed that a "shape" is what is left when the "nuisance" degrees of freedom corresponding to translation, scaling, and rotation (i.e., pose) have been filtered away, as was proposed by Kendall [8]. Underlying this definition is the idea that the shape itself is an object with an infinite number of degrees of freedom, such as a continuous curve or surface. The analysis of shape is then the

General Shape Analysis:
- Define a suitable representation for our shapes, hence a space of shapes.
- Define an energy or distance function on that space, a measure of the similarity between any two such shapes.

Pairwise Shape Analysis:
- To provide a continuous, optimal path, that allows interpolation between the two shapes.
- To give a quantitative measure of the degree of similarity.

Groupwise Shape Analysis:
- To provide a method of interpolation across a training set of shapes.
- To provide a means of analysing the statistics of the training set of shapes.

Fig. 1. The key aspects of general, pairwise, and groupwise shape analysis.

related problems of representing the infinite dimensional curves, defining metric distances between different shapes, finding optimal correspondences between shapes, and performing statistical analysis on the manifold of shapes. In Fig. 1 we identify these problems within the general framework, and then in the concrete cases of pairwise and groupwise shape analysis. The last of these, which has been generally neglected, is probably the most important for real shape analysis.

In oder to remove the transformation group of rotation, translation, scaling, a method of shape alignment (such as Procrustes analysis) is often defined *a priori*. In fact, the same approach has been taken to defining a correspondence between shapes as well, and this causes problems. Defining a correspondence between shapes makes assumptions about how to best interpolate between shapes. When such an *a priori* choice is applied to interpolation between a pair of shapes, the result may not reflect what is actually seen when we come to consider a group of shapes. This is shown in Fig. 2, where two choices of correspondence lead to different interpolated shapes. We contend that the issue of correspondence should be left open, and the various alternative hypotheses considered, so that the result of interpolation may then be chosen to agree with what is seen from the entire group of shapes.

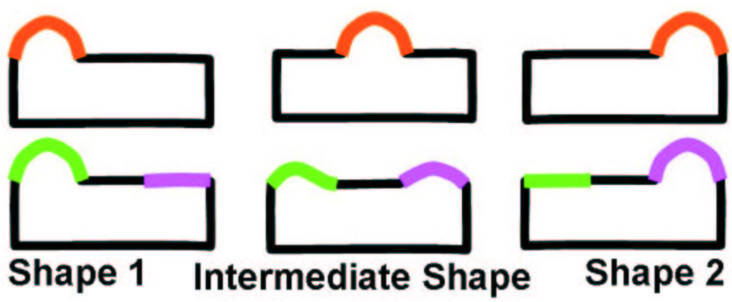

Fig. 2. For the same pair of shapes (left & right), different correspondences (indicated by colours) leads to different hypotheses as to the intermediate shapes.

In this paper we present a framework for shape analysis, and place many of the recent papers on the topic within that framework. We will first show how the question of shape correspondence is inextricably linked to the use of a Riemannian metric, both in terms of the tangent space, and in terms of the associated Levi-Civita connection, and then discuss how to optimise correspondence, both pairwise and groupwise.

2 The Metric

We start by considering parameterised shapes, where a shape is described by a function $c(\theta) \in \mathbb{R}^n$, where $\theta = \{\theta_\alpha : \alpha = 1, \dots m\}$ represents a point in the m-dimensional parameter space M. All that we require is that there is a continuous, one-to-one correspondence between the shape and the parameter space, so that $\{c_{\theta_\alpha} \doteq \frac{\partial c}{\partial \theta_\alpha}\}$ are not all zero (i.e., no pleats). The mapping from the parameter space to $\mathbb{R}^n$ can be either a smooth immersion (self-intersections in $\mathbb{R}^n$ permitted) or a smooth embedding (self-intersections in $\mathbb{R}^n$ *not* permitted). In order to define a Riemannian metric between shapes, we need to consider the tangent space to our space of parameterised shapes at the shape c, the space of all possible infinitesimal deformations of $c(\theta)$. This can be represented by:

$$c(\theta) \Rightarrow c'(\theta) \doteq c(\theta) + \epsilon h(\theta), \;\; \epsilon \ll 1, \tag{1}$$

where $h(\theta)$ is a continuous and (piecewise) smooth $\mathbb{R}^n$-valued vector field[3], defined everywhere on M. This hence represents the infinitesimal deformation of our original shape in the direction defined by $h(\theta)$. This gives the most general description possible of deforming a shape, in that the deformation of every point on the original shape is given, with the only constraint being that the deformation is smooth and continuous. However, this very generality means that translations (constant vector fields), rotations, scalings, and re-parameterisations of the original curve (purely tangential vector fields) are also included as elements of the space of vector fields.

It is important to note that this definition defines a point-to-point *correspondence* between two infinitesimally-separated shapes c and c', given by the value of the parameter, so that $c(\theta)$ corresponds to $c'(\theta)$. This mapping can also always be made one-to-one, because even if, for some value of ϵ, there are points where $c'_{\theta_\alpha} = c_{\theta_\alpha} + \epsilon h_{\theta_\alpha} = 0 \;\forall\; \alpha$, we can always adjust ϵ so that it is not the case, since there is always *some* α such that $c_{\theta_\alpha} \neq 0$.

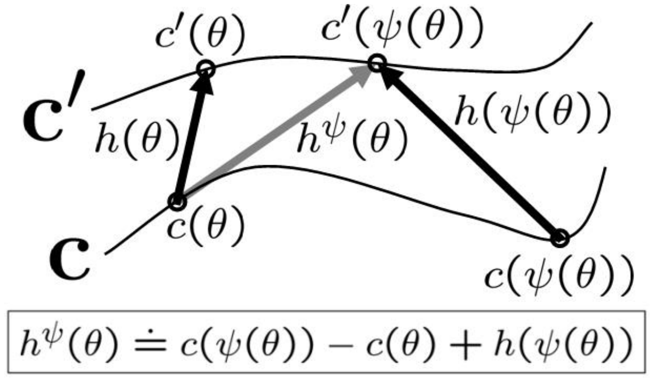

Fig. 3. Correspondence between two curves, c and c', showing how the vector field changes $h \to h^\psi$ as c' is reparameterised.

---

[3] In what follows, we will often absorb the factor of ϵ into the definition of h, hence take $c(\theta) + h(\theta)$ as our deformed shape.

We can also change the correspondence between the infinitesimally-close shapes, by reparameterising the curve c': $c'(\theta) \rightarrow c'(\psi(\theta))$ (see Fig. 3). Hence changing the correspondence between the two shapes means changing the vector field $h \rightarrow h^{\psi}$, whilst keeping the set of points $\{c'(\theta) : \theta \in M\}$ fixed.

A Riemannian metric on our space of shapes is then an assignment of an inner product between elements of the tangent space at c, denoted by $G_c(h,k)$. The *energy* and *length* of the path segment between the (infinitesimally) separated shapes $c(\theta)$ and $c(\theta) + h(\theta)$ are then given by:

$$\delta\text{Energy} = G_c(h,h), \;\; \delta\text{Length} = (G_c(h,h))^{\frac{1}{2}} . \tag{2}$$

We take as an exemplar (open or closed) parametric planar shapes (PPS). Defining our basic notation:

Curves and vectors in the Argand plane: $c(\theta), h(\theta) \in \mathbb{R}^2 = \mathbb{C}$.

Arc-length: $(ds)^2 \equiv |c(\theta + d\theta) - c(\theta)|^2$.

$M =$ unit line or circle for open or closed curves, $\theta \in [0,1]$ or: $\theta \in [0, 2\pi]$.

Derivatives: $D_s \equiv \dfrac{\partial}{\partial s}, \;\; D_\theta = \dfrac{\partial}{\partial \theta}, \;\; D_\theta f(\theta) \equiv f_\theta$.

Unit tangent vector: $v_c = c_s = \dfrac{c_\theta}{|c_\theta|}$, Unit normal to curve: $n_c = i.v_c$.

Following Michor & Mumford [12, 11], and Younes et al. [14] we can then define the local and almost-local metrics:

$$\text{Local:} \quad G_c(h,k) = \int \langle Lh(\theta), Lk(\theta) \rangle d\mu(\theta), \tag{3a}$$

$$\text{Almost local:} \quad G_c(h,k) = \int \Phi_c(\theta) \langle Lh(\theta), Lk(\theta) \rangle d\mu(\theta), \tag{3b}$$

where $d\mu(\theta)$ is our integration measure (which may involve the total length of the curve l_c), $\Phi_c(\theta)$ is a function of l_c and the curvature κ_c, (where $\kappa_c n_c \doteq D_s v_c(s)$), and L is some differential operator, such as $(D_s)^n$. The inner product $\langle \cdot, \cdot \rangle$ will usually be the Euclidean inner product on $\mathbb{R}^2$.

The shortest-path $c(\theta, t)$ (the *geodesic*) between two finitely-separated shapes $c(\theta, 0)$ and $c(\theta, 1)$ is then given by minimising the total energy[4]:

$$\int_0^1 G_{c(\theta,t)}(c_t, c_t) dt, \tag{4}$$

wrt variations of the path $c(\theta, t)$, whilst keeping the endpoints fixed. If we take an algebraic approach, we will need to compute the variation of terms such as the metric, as we vary the path $c(\theta, t)$. It is convenient to define the variational, Gâteaux derivative:

$$D_{c,m} G_c(h,k) \doteq \left. \frac{d}{d\lambda} \right|_{\lambda=0} G_{c+\lambda m}(h,k), \tag{5}$$

[4] As in the case of finite-dimensional Riemannian geometry, the total length and total energy have the same minimizer, and the total energy is simpler to work with.

where $\lambda m(\theta)$ represents an infinitesimal variation of the curve $c(\theta)$ in the direction $m(\theta)$. We can then define derivatives of the metric H_c and K_c, where:

$$G_c\left(K_c(m,h),k\right) \doteq D_{c,m}G_c(h,k). \;\; G_c\left(m,H_c(h,k)\right) \doteq D_{c,m}G_c(h,k). \quad (6)$$

The formal solution is then given by the geodesic equation:

$$c_{tt} = \frac{1}{2}\left(H_c(c_t,c_t) - 2K_c(c_t,c_t)\right). \quad (7)$$

The path ahead in the pairwise case now seems simple – we make our choice of metric, solve the geodesic equation (either numerically [12] or algebraically if possible [14]), and hence compute the geodesic between any pair of shapes. Given the geodesic path between the shapes, it is then trivial to compute the length, hence assess their level of shape similarity.

However, we have forgotten the issue of correspondence. As noted previously (see Fig. 3), even for the same sequence of physical shapes, altering the correspondence alters the vector fields, and hence will in general give a different distance between shapes.

2.1 Global & Local Re-Parameterisations

Consider a general path on the space of curves $c(\theta,t)$, and suppose that we now consider a *global* re-parameterisation of the path, so that $c(\theta,t) \to c(\psi(\theta),t)$. This will leave the correspondence unchanged along the path, and the vector fields unchanged, but will change the integration measure, and hence the path energy, unless we choose a metric that is *invariant* under re-parameterisation. That is, $G_c(h,k)$ is invariant under the transformation $\theta \to \psi(\theta)$. This invariance is important, otherwise we can reduce the energy cost of finite deformations to zero by simply identifying a suitable choice of re-parameterisation (see [11], §3.1).

However, even for such a re-parameterisation invariant metric, the path energy will not be invariant under a *local* re-parameterisation, $\psi(\theta,t)$. Since the re-parameterisation function is different at different points along the path, it changes the correspondence and the vector fields c_t (see Fig. 3), and thus in general, it changes the path energy. Hence using re-parameterisation-invariant metrics does not solve the pairwise correspondence problem.

It would seem a sensible course to nail the issue of correspondence down at the start, and this is the approach taken by Michor & Mumford. Their approach rests on the observation that a *purely* tangential vector field generates not changes of shape of a curve c, but changes of parameterisation of c. A general vector field can be decomposed into tangential and normal components, where:

$$h(\theta) = h^{\perp}(\theta) + h^{\|}(\theta), \;\; h^{\|}(\theta) \doteq \langle h(\theta), v_c(\theta)\rangle v_c(\theta), \;\; h^{\perp}(\theta) \doteq \langle h(\theta), n_c(\theta)\rangle n_c(\theta),.$$

This means that purely normal vector fields are perpendicular to the directions in shape space that represent pure re-parameterisations, hence always perpendicular to the orbit of a shape under the action of the diffeomorphism group of the parameter space M. They then use what we will call the *normality prescription*, where they project out the degrees of freedom in the vector fields that correspond to re-parameterisation, replacing the inner product in $\mathbb{R}^n$ with a term that depends only on the normal components. For the H^0 metric:

$$\langle h, k\rangle \rightarrow \langle h^{\perp}, k^{\perp}\rangle \equiv \langle h, n_c\rangle\langle k, n_c\rangle$$

This normality prescription is not without some undesirable effects. In [11], Michor and Mumford state that they were looking for the simplest Riemannian metric, the obvious candidate being the re-parameterisation-invariant H^0 metric:

$$G^0_c(h,k) \doteq \int \langle h(\theta), k(\theta)\rangle |c_\theta| d\theta \equiv \int \langle h(s), k(s)\rangle ds,$$

which can be re-written in terms of the arc-length parameter s and which is obviously related to the simple sum-of-squared-distances metric for polygonal shapes represented by a finite number of points.

However, if we apply the normality prescription, then this metric goes horribly wrong (see [11], §3.10), and all geodesic distances can be reduced to zero! It was this problem with the H^0 metric that led Michor & Mumford to consider higher-order derivative terms (local metric (3a)), or the addition of curvature-dependent terms (almost-local metric (3b)), whilst still keeping the normality prescription.

What the normality prescription does, in effect, is assign an equal cost to all possible correspondences, by ignoring the tangential component that generates such changes of correspondence. This is not the same as choosing one particular way of assigning correspondence, which would be equivalent to assigning an infinite cost to all other correspondences. This freedom to move *along* shapes is one way of seeing why the geodesic distances can be reduced to zero − points can move from one shape to the second for zero cost if they can complete the journey by sliding along a shape, and adding saw-teeth stretching between the two shapes, as intermediate shapes are allowed this under the prescription.

Our contention is that since the normality prescription does not work for the H^0 metric, then it seems to us unwise to continue with it.

But perhaps the correspondence problem can be improved by moving to a different representation of shape? We mention here two alternatives.

The first is the elastic approach [6, 7, 15], which rather than the shape function $c(s)$, uses the speed function c_s. In particular, the square-root-elastic (SRE) representation uses the variable and metric:

$$q(\theta) \doteq \frac{c_\theta(\theta)}{\sqrt{|c_\theta(\theta)|}},\;\; G^{SRE}(\Delta q_1, \Delta q_2) \doteq \int \langle \Delta q_1, \Delta q_2\rangle d\theta, \tag{8}$$

where $\Delta q_1, \Delta q_2$ are elements of the tangent space to the space of speed functions. However, we can re-write this in terms of the elements of the tangent space to the space of shape functions, to find that:

$$G^{SRE} \Rightarrow \int \frac{1}{|c_\theta|}\left(\langle h_\theta, k_\theta\rangle - \frac{3}{4}\langle h_\theta, v_c\rangle\langle k_\theta, v_c\rangle\right) d\theta.$$

Hence we see that the SRE metric is just a particular combination of re-parameterisation invariant terms involving just the first derivatives of the vector field, chosen so as to give the metric a simple form in the space of speed functions. The issue of correspondence is as in the case of parameterised shape, and we note that in [7], Joshi et al. find the detailed correspondence and the path between a pair of shapes by explicit optimisation.

The second, and more intriguing approach, is that given by the use of conformal mappings [13]. For any simple closed planar shape, there always exists a conformal (angle-preserving) mapping from the interior of the shape to the unit disc. This hence defines a *conformal parameterisation* of the shape, in terms of the mapping between points on the curve and points on the unit circle. We can also consider the inverted shape (for a shape $c \subset \mathbb{C}$, and a point z_0 in the interior of the shape, the inverted shape is given by $\frac{1}{c-z_0}$), and find the conformal parameterisation of that. In general, the conformal parameterisation of the shape and the inverted shape are different, and the difference between the two is a diffeomorphism of the unit circle. Sharon & Mumford [13] show that it is possible to (almost) uniquely reconstruct a shape purely from knowledge of this element of the diffeomorphism group of the unit circle, and hence establish a way to represent shapes in terms of this diffeomorphism group. They can then apply metrics on the diffeomorphism group to generate a metric on the space of shapes, and also, from the group multiplication, obtain an intriguing multiplication of two shapes to give a third shape. This method has been generalized by Lui et al. [9] to the case of planar objects with other topologies. However, the entire method still rests on favouring a particular method of parameterisation over any other, and is limited to purely planar shapes.

Rather than trying to consider each method of shape representation and each method of defining a distance between shapes on a case-by-case basis, we will instead move on to consider the general geometric setting for shape distances.

3 Correspondence & Connections

In this section, we first start by discussing what it is about shapes and spaces of shapes that makes them distinct from other spaces. In particular, we will focus on the notion of spatial localization.

Let us suppose we have some general method of shape representation and a shape space S, shapes $c \in S$, and a Riemannian metric $(G_c(\cdot,\cdot))$ defined on such a space. By general, we mean that our shape representation should be such that it can also represent infinitesimal, localised deformations of any permitted shape – this can be thought of in terms of growing an infinitesimal bump or pit at any point on any shape. For a finite-dimensional representation of shape, such as the simple polygonal or spline-based representations, this requirement becomes the ability to move only a single point, plus the ability to increase the number of points used in the representation as required. We will then associate such bumps and pits with localized elements of the tangent space to the space of shapes, which we will refer to as *bump vectors*. For a general shape, growing a bump at one of two distinct points on the shape should be recognized as distinct directions in the tangent space, since they generate distinct shapes in the finite limit. Hence we will refer to a shape c, where A and B are distinct points on the shape, and distinct elements of the tangent space at c, $k_A \in T_cS$ and $k_B \in T_cS$, which correspond to growing a localized bump at point A or at point B.

It is important to note that our intuitive ideas of shape and shape change rest on the notion of locality. In particular, we have the idea that, in general, spatially-separated small perturbations of a single shape represent distinct degrees of freedom, provided these perturbations are sufficiently far apart. In terms of the metric and our bump vectors, this can usefully be stated in the form that $G_c(k_A, k_B) \to 0$ as $|A - B|_{\mathbb{R}^n}$ increases[5]. Hence we restrict ourselves to shape metrics that have *some* notion of spatial locality and localization. The local and almost-local metrics in (3a) & (3b) obviously have this property, and we will consider a more non-local metric later in this section.

For a general space, the tangent spaces at two distinct points are not equivalent, just as the tangent plane to a sphere at the pole is a different plane to the tangent plane at a point on the equator. A *connection* provides a recipe (called *parallel transport*) for mapping elements of the tangent space at one point into elements of the tangent space at any other point. For the particular case of a Riemannian metric, there is a unique (torsion-free) connection (the Levi-Civita connection) that preserves the metric. In terms of the Gâteaux derivatives of the metric we defined earlier (6), this connection is given by:

$$\mathit{\Gamma}_c(h,k) \doteq \frac{1}{2}\left(H_c(h,k) - K_c(h,k) - K_c(k,h)\right), \tag{9}$$

where it should be noted that this formula is general, and not specific to the case of a parametric representation of shape. An element k of the tangent space at the point c can then be parallel-transported by an infinitesimal amount ϵ in the direction h, to give the element of the new tangent space, which can be written as:

$$k \in T_cS \;\;\to\;\; k + \epsilon \mathit{\Gamma}_c(h,k) \in T_{c+\epsilon h}S. \tag{10}$$

Tangent space vectors can then be parallel-transported a finite distance along paths in the space by integrating up the above result, and in general, the exact result will depend on the path chosen, even with fixed endpoints[6].

We can now see how this construction of parallel-transport applies to the issue of correspondence between shapes. If we take a bump vector k_A on one shape, we can parallel-transport this tangent-space vector along a path between shapes, and hence generate the corresponding element of the tangent space on our second shape. If this new element is also localized, then its location provides us with a (rough) correspondence between the shapes.

We take as our example a metric on parametric shapes, where the connection can be computed in closed form. In this case, we already know the answer

[5] That is, they become *orthogonal* at sufficient spatial separation.

[6] This dependence of the result of parallel transport on the exact path taken is one definition of the *curvature* of the underlying manifold.

as to the correspondence between shapes we expect to recover, it is just the correspondence given by parameter value. We take a translation-invariant H^1 metric:

$$G_c(h,k) = \int \frac{1}{|c_\theta|}\langle h_\theta, k_\theta\rangle d\theta,$$

which is also re-parameterisation invariant. Unlike the H^1 metric used by Younes [15], it does not include the extra factor of $\frac{1}{l_c}$, which would make it scale-invariant. To compute the Gâteaux derivative (5), we note that the only piece that varies is the $|c_\theta|$ term, which gives:

$$D_{c,m}G_c(h,k) = -\int \frac{1}{|c_\theta|^2}\langle v_c, m_\theta\rangle\langle h_\theta, k_\theta\rangle d\theta.$$

To compute the connection, we take a specific form for the bump vector $k(\theta)$ (a top-hat function, given by a constant vector α between θ_0 and θ_1, and zero elsewhere), so that:

$$k_\theta(\theta) = a\left(\delta(\theta-\theta_0) - \delta(\theta-\theta_1)\right).$$

If we then also let $\theta_1 \to \theta_0$ (so that terms such as $f(\theta_1) - f(\theta_0)$ can be taken to vanish in the limit), then using the definitions (6) & (9), we find that

$$\Gamma_c(h,k)(\theta) = \frac{1}{2|c_\theta|}\left[\langle v_c, h_\theta\rangle\alpha + \langle v_c, \alpha\rangle h_\theta - \langle\alpha, h_\theta\rangle v_c\right](\theta) \text{ if } \theta = \theta_0, \text{ else } 0.$$

Hence, as we might have expected, any change in $k(\theta)$ under parallel transport is localized at θ_0. Note also that if the change of shape $h(\theta)$ is locally a translation ($h_\theta(\theta_0) = 0$), then there is no change in $k(\theta)$ under parallel transport, which reflects the fact that the metric is an H^1 metric. And in general, since the result contains terms in the three directions α, $v_c(\theta_0)$, and $h_\theta(\theta_0)$, the direction of $k(\theta)$ may change under the transport, even though the foot-point remains unchanged.

Finally, we consider an extension to the metrics we have considered so far. In [5], Glaunès et al. considered a non-local curve-matching energy term for finitely-separated curves (derived from a norm on the space of currents), which was incorporated within the large deformation diffeomorphic mapping framework. This energy was of the form:

$$E(c,c') = F(c,c) - 2F(c,c') + F(c',c'), \tag{11a}$$

$$F(c,c') \doteq \int d\theta \int d\phi\ \mathcal{K}(c(\theta), c'(\phi))\langle c_\theta, c'_\phi\rangle, \tag{11b}$$

where $c(\theta)$ and $c'(\phi)$ are two parametric curves, and $\mathcal{K}(x,y) \equiv \mathcal{K}(|x-y|)$ is a kernel function (such as a Gaussian). If we take an infinitesimal difference of curves, $c' = c + \epsilon h$, and make the approximation that:

$$\mathcal{K}(c(\theta) + \epsilon h(\theta), c(\phi) + \epsilon h(\phi)) \approx \mathcal{K}(c(\theta), c(\phi)),$$

then we obtain the final non-local Riemannian metric in the form:

$$G_c(h,k) = \int d\theta \int d\phi\ \mathcal{K}(c(\theta), c(\phi))\langle h_\theta(\theta), k_\phi(\phi)\rangle.$$

This is translation and re-parameterisation invariant, and an obvious generalization of the H^1 metric that we considered previously (and a similar generalization can obviously be applied to the other metrics considered earlier (3a) & (3b)).

Note that the original energy term for finitely-separated curves *does not* involve an explicit correspondence between the curves based on parameter value, and was in fact invariant to local re-parameterisations. However, when we made the simplifying assumption to replace $\mathcal{K}(c, c')$ etc. by $\mathcal{K}(c, c)$, we removed this invariance, and instead replaced it by the correspondence according to parameter that we had in the cases of the local and almost-local metrics. It should seem that we have taken a case without correspondence, and put it back in by hand!

This is not quite the case: the original formulation assigned an *explicit* correspondence based on points in the plane. The positioning of the curves in the plane then allowed the distance between points on the two curves to act to establish the notion of locality and the meaning of local differences in shape between the curves. The optimisation over diffeomorphisms *of the plane* that Glaunès et al. [5] then use to match curves is the equivalent of the optimisation over correspondence that we propose. We note that other formulations (such as shape representation using distance maps), also employ point-to-point correspondence across the plane as an alternative to point-to-point correspondence between shapes.

4 Optimising Correspondence

We have seen from the previous section that the question of correspondence is inextricably tied up with the use of Riemannian metrics on shape spaces. There are then three possible approaches to dealing with this issue:

(1) Define a method of determining correspondence *a priori*. Examples would be basing correspondence on equal fractional arc-length, or the use of the conformal parameterisation that formed part of the work in [13]. The problems are that this choice is essentially arbitrary, and that a method that gives sensible interpolation for pairs of shapes from one class may not give suitable results for shapes from a different class.

(2) Try to factor-out these degrees of freedom. This is essentially the approach taken in the normality prescription case (see §2.1), where all possible correspondences are assigned an equal weight. But as we have already noted, the simplest H^0 metric fails in this case, which does not seem a desirable result.

(3) Determine the optimum correspondence in a data-driven fashion. This can then obviously be extended to find the optimum pose.

In the pairwise case, given the absence of any other data, the only information we have that distinguishes between different correspondences is the geodesic distance itself. Hence, it would seem sensible to allow this to define the optimum pairwise correspondence, despite the complication of a further optimisation step. This was the approach taken by Joshi et al. [7], in the case of the square-root-elastic metric.

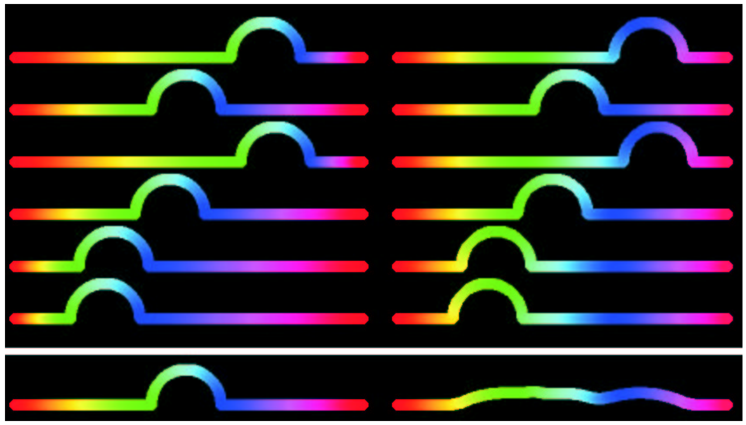

Fig. 4. A set of training shapes, correspondence indicated by colour, for two different correspondences (*Left:* correct, *Right:* arc-length). *Bottom:* The Euclidean mean.

When we come to the groupwise case, why can we not just take the pairwise correspondences defined as above to create a groupwise correspondence? The problem is that, in general, the correspondence defined between shapes A and B, and between A and C, will not agree with that defined between B and C. Geometrically, this is because parallel-transport around a closed loop gives a result other than the identity (which is a definition of *curvature*). This question did not arise in earlier work (such as Davies et al. [3]), since the Riemannian metric used there was Euclidean.

The obvious solution is to define correspondence via some reference shape, which makes the groupwise correspondence consistent by construction, the obvious candidate being the Karcher mean. The mean shape also gives us another advantage, in that we can use length/area on the mean in order to define our integration measure (3a), as was done in [3].

This now gives us our basic framework for groupwise shape analysis. However, we still have to define the objective function that we are going to use to define the optimum correspondence. The simplest suggestion is to just repeat the procedure we used in the pairwise case, and take the sum of geodesic distances to the mean (the *compactness*) to define both the mean (for fixed correspondence), and the optimum correspondence. However, this repeated-pairwise approach does not always work.

To give an example, consider the set of curves shown in Fig. 4, where we have used the Euclidean metric, and integration measure computed on the mean shape [3]. We take two different correspondences, the correct correspondence (on the left), and arc-length correspondence (on the right). The correct correspondence gives the mean as just another bump, whereas the arc-length case gives a shape unlike any seen in the training set. However, if we follow the same colour across examples, we see that distances from the mean will be larger for the correct correspondence, since the straight-line portions of the curve have to stretch

	Correct	Arc-Length
	Mean (Spread)	Mean (Spread)
Eucl.	2.03 (1.37)	0.68 (0.11)
SRE	2.10 (1.39)	1.05 (0.25)

Table 1. Mean and spread of distances to mean shape for two choices of correspondence, and using two different metrics.

and compress to accommodate the motion of the bump, as well as the motion associated with the bump itself. In the arc-length case, although the mean is not a bump, movements are minimal, hence this correspondence will be measured as being more compact. For the same group of shapes, we also repeated the analysis using the speed function representation and the SRE metric, as in (8). The distance to their respective means, and the spread of values, are given in the Table. It can be seen that, for both metrics, compactness fails as a means of choosing the correspondence that is in accord with the mode of variation seen in the input shape data.

5 Discussion

In this paper we have provided a framework for shape analysis and discussed why the related problems of metric, correspondence, and connection all need to be selected in a data-dependent way. This is particularly clear from our last example. It could be argued that in this case a metric that gave greater weight to shape similarity based on curvature would give a better result. However, that would miss the essential point: what distinguishes the correct correspondence in this case is not curvature per se, but the commonality of structure across the group of shapes. The association of the edges of the bumps with regions of high curvature should then be seen as an accident of the artificial shape construction. From the point of view of *modelling*, it is obviously desirable that the mean should reflect the common structure seen across the group, and unless we have correctly identified the common structure across examples, we will be unable to correctly represent the variation of this structure.

We note that the groupwise case is more complicated than the pairwise case, in that we don't want to just interpolate between pairs of example, but across the whole sub-space in which the training data lies. It was for this reason that more sophisticated groupwise objective functions (such as MDL [3]) were introduced for the case of Euclidean shape spaces. The construction of such objective functions will be more complicated in the non-flat case, since we can no longer construct the simple pdf models on the shape space. Given this, we might ask why we might need to use metrics other than the Euclidean one? One illustrative example is where parts of an object undergo motion which is a rotation (such as the thumb of a hand, see [2], Fig. 9.9). The MDL correspondence in this case tends to linearize the motion by allowing points on the tip to slide, which is not quite the correct correspondence from a physical point of view. We will be considering this further in the future.

We note that there do exist methods for modelling on non-flat shape spaces, and these entail constructing models on the tangent space at the mean (e.g., principal geodesic analysis [4]), which again shows that it is necessary that the mean itself is similar to the shapes seen in the group. Developing alternative objective functions to compactness, in the spirit of MDL, is the obvious next step, but is beyond the scope of the current paper.

In §3 we identified that the method of Glaunès et al. used a rotationally invariant kernel $\mathcal{K}(x, y)$. This kernel defines an inner product between parametric curves $c(\theta)$ and $c'(\phi)$ [10], which clearly induces a particular Riemannian metric on the space. This has been considered in the area of machine learning, where the kernel mapping of a Support Vector Machine performs essentially the same mapping. There, Burges [1] looked for locally invariant kernels under some symmetry and identified how the induced metric can be expressed in closed form. In future work we will follow up this line to identify whether it is possible to choose the kernel in a data-driven way for groupwise shape analysis.

Acknowledgements: Our thanks to S. H. Joshi for making available his MATLAB implementation of the SRE metric.

References

1. Burges, C.: Geometry and invariance in kernel based methods. In: Schölkopf, B., Burges, C., Smola, A. (eds.) Advances in Kernel Methods – Support Vector Learning, pp. 89–116. MIT Press (1999)
2. Davies, R., Twining, C.J., Taylor, C.J.: Statistical models of shape: optimisation and evaluation. Springer (2008)
3. Davies, R.H., Twining, C.J., Cootes, T.F., Taylor, C.J.: Building 3-D statistical shape models by direct optimization. IEEE Transactions on Medical Imaging 29(4), 961–81 (2010)
4. Fletcher, P.T., Lu, C., Pizer, S.M., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. IEEE Transactions on Medical Imaging 23(8), 995–1005 (2004)
5. Glaunès, J., Qiu, A., Miller, M.I., Younes, L.: Large deformation diffeomorphic metric curve mapping. International Journal of Computer Vision 80(3), 317–336 (2008)
6. Joshi, S.H., Klassen, E., Srivastava, A., Jermyn, I.: A novel representation for riemannian analysis of elastic curves in $\mathbb{R}^n$. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1643–1649 (2007)
7. Joshi, S.H., Klassen, E., Srivastava, A., Jermyn, I.: Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves. In: Proceedings of EMMCVPR. LNCS, vol. 4679, pp. 387–398 (2007)
8. Kendall, D.G.: The diffusion of shape. Advances in Applied Probability 9(3), 428–430 (1977)
9. Liu, L.M., Zeng, W., Yau, S.T., Gu, X.: Shape analysis of planar objects with arbitrary topologies using conformal geometry. In: Proceedings of ECCV 2010. LNCS, vol. 6315, pp. 672–686 (2010)

10. McLachlan, R.I., Marsland, S.: N-particle dynamics of the Euler equations for planar diffeomorphisms. Dynamical Systems 22(3), 269 – 290 (2007), http://www-ist.massey.ac.nz/smarsland/PUBS/DynSys07.pdf
11. Michor, P.W., Mumford, D.: Riemannian geometries on spaces of plane curves. Journal of the European Mathematical Society 8, 1–48 (2006)
12. Michor, P.W., Mumford, D.: An overview of the riemannian metrics on spaces of curves using the hamiltonian approach. Applied and Computational Harmonic Analysis 23, 74–113 (2007)
13. Sharon, E., Mumford, D.: 2d-shape analysis using conformal mapping. International Journal of Computer Vision 70(1), 55–75 (2006)
14. Younes, L., Michor, P.W., Shah, J., Mumford, D.: A metric on shape space with explicit geodesics. Rendiconti Lincei - Matematica e Applicazioni 9, 25–37 (2008)
15. Younes, L.: Computable elastic distances between shapes. SIAM Journal of Applied Mathematics 58(2), 565–586 (1998)