SURVEYING SHAPE SPACES

CHARLESS C. FOWLKES

"There are manifoldnesses in which the determination of position requires not a finite number, but either an endless series or a continuous manifoldness of determinations of quantity. Such manifoldnesses are, for example, the possible determinations of a function for a given region, the possible shapes of a solid figure, and so on"
—Bernhard Riemann (as translated by W.K. Clifford)

1. Introduction

One doesn't have to look far to find interesting examples of manifolds. A glance around the room usually reveals our favorite examples: a plane, a cylinder, with a little luck, a donut. Riemannian geometry allows a precise classification of the "shape" of these surfaces as well as far more abstract objects. Backing away from the hard constraints of isometry, we can ask: what surfaces are "nearly isometric" in the sense that they are isometric after a small, smooth deformation. If we can establish a smooth manifold structure on a space of possible shapes, then we are free to play with all sorts of fun applications such as finding the shortest path between two shapes (morph the president into a banana) or studying the statistics of a collection of shapes (find out what the average person on the street actually looks like).

This paper provides an overview of results on some simple shape spaces and will hopefully convince the reader that this area is rich with geometrical structures. Section 2 presents results for the space of arrangements of k points in $\mathbb{R}^n$ modulo translation, rotation and scale. In Section 3 we examine three related descriptions of the space of continuous curves in the plane.

2. The shapes of finite point sets

D. G. Kendall[7] pioneered the study of shape for labeled point sets. He proposed to study shape as that which "is left when the effects associated with translation, scaling and rotation are filtered away". Kendall's applications are in the field of statistics, for example, modeling the locations of artifacts found at an archaeological site. T.K. Carne [3] proposed the following more general definition. Let G be a Lie group that acts smoothly on M. The *shape space* $\Sigma^k(M, G)$ is defined to be the set of orbits of the action of G on the k-fold product M^k. Aspects of this sort of shape space have also been studied elsewhere i.e. mechanics [12]. Our presentation will focus on the case where $M = \mathbb{R}^m$ and G is the set of similarity transformations. Further reading is best provided by [15] which gives a comfortable introduction with statistical applications or by [3] which supplies a full dose of generality.

Consider a finite set of points represented by the matrix $X = (x_1, x_2, \ldots, x_k) \in \mathbb{R}^{m\times k}$ where columns give the coordinates of each point in $x_j \in \mathbb{R}^m$. Translation and scale are easily dealt with by utilizing a coordinate system whose origin lies at the center of mass and whose axes are scaled so that the sum of squared distances to our points from the origin is 1. Normalizing for scale is equivalent to dividing by $r = \|X\| = \sqrt{Tr(XX^T)}$ [1]. This space, given coordinates by X is referred to as the *pre-shape space*. It is easily seen to be the unit sphere $\mathbb{S}^{m(k-1)-1}$. Rotation is "filtered away" by viewing the pre-shape space modulo the rotation group $\mathbf{SO}(m)$ where the action of $\mathbf{SO}(m)$ is from the left on the columns of X, rotating the configuration of points about the origin. This space of orbits is the *Euclidean shape space* $\Sigma_m^k \equiv \Sigma(\mathbb{R}^m, \mathbf{SE}(m)) \simeq \mathbb{S}^{m(k-1)-1}/\mathbf{SO}(m)$. Let $\pi : \mathbb{S}^{m(k-1)-1} \to \Sigma_m^k$ be the map which takes a pre-shape X to its equivalence class.

For $m = 1, 2$, $\mathbf{SO}(m)$ acts freely on the elements of our pre-shape space so that Σ_m^k is a smooth manifold (a point when $k < 3$). When $m \geq 3$, $\mathbf{SO}(m)$ acts freely except at a set of singular points where the rank of X is less than $m - 1$. For example, there is some rotation that leaves fixed any set of co-linear points in $\mathbb{R}^3$. Let $\mathcal{D}_{m-2}$ be the set of pre-shape matrices with rank less that $m - 1$. Away from $\pi(\mathcal{D}_{m-2})$, the image of these singularities, the quotient map π is a submersion. The topology of Σ_m^k is defined so that U is open if $\pi^{-1}(U)$ is, making Σ_m^k a continuous image of the pre-shape space sphere and hence compact. The spaces Σ_m^k are also arc-wise connected except for $\Sigma_1^2 \simeq \mathbb{S}^0$. $\Sigma_1^3 \simeq \mathbb{S}^1$ is the only arc-wise connected Euclidean shape space that isn't simply connected [7].

2.1. The Riemannian geometry of Σ_m^k.

The basic tool for understanding the geometry of Σ_m^k is to define a metric on the shape space so that π is a Riemannian submersion (on the complement of $\mathcal{D}_{m-2}$). We can proceed locally by letting the non-singular part of Σ_m^k inherit the inner product $\langle X, Y\rangle = Tr(XY^T)$ from the pre-shape space. Alternately, the usual round metric on the pre-shape space descends to a well defined metric on Σ_m^k called *Procrustes distance*[2]. For two shapes $\pi(X), \pi(Y) \in \Sigma_m^k$ this distance can be written as

$$d(\pi(X), \pi(Y)) = \inf_{Z,W} \{\cos^{-1}(\langle Z, W\rangle) : \pi(Z) = \pi(X), \pi(W) = \pi(Y)\}$$

Because of the symmetry of the pre-shape space, we only need to take the infimum over one of the fibers

$$d(\pi(X), \pi(Y)) = \inf_{Z} \{\cos^{-1}(\langle Z, Y\rangle) : \pi(Z) = \pi(X)\} \tag{2.1}$$

These global and local definitions are in fact identical.

Proposition 2.1. The Riemannian metric on the regular part of Σ_m^k determines a global distance which is identical to Procrustes distance. See [3] or [7] for details.

Since π is a Riemannian submersion when restricted to regular points, we can express geometric invariants of the shape space in terms of those of the pre-shape

[1] We will exclude from further consideration the completely degenerate case where all points are at the same location (also exclude the case $k = 1$) so $r > 0$. As Kendall puts it, "The totally degenerate k-ads are omitted because from one point of view they have no shape, and from another they 'almost' have every shape"

[2] Procrustes was a gruesome character from Greek legend who invited passing strangers into his home along the road to Athens. He offered them a night's rest in a magical bed which he claimed had the property of exactly matching whomsoever lay down upon it. As the unwary traveler soon found, this miracle was accomplished either by stretching on the rack or chopping off the legs. Procrustes, whose name means "he who stretches" was eventually lain to rest in his own bed by Theseus

space. The tools of choice are O'Neill's results on Riemannian submersions[13]. For example

Theorem 2.2 (O'Neill 1966)**.** *Suppose that M is submersed into $\tilde{M}$ and $\tilde{X}, \tilde{Y}$ are the horizontal lifts of X and Y; then the sectional curvatures are related by:*

$$\tilde{K}(\tilde{X}, \tilde{Y}) = K(X, Y) - \frac{3}{4}\|[\tilde{X}, \tilde{Y}]^v\|^2$$

Proof. The appendix gives a proof of the relation between the Riemannian curvatures R and $\tilde{R}$. The expression for sectional curvature follows readily. □

Since the pre-shape space is a sphere with a constant sectional curvature of 1, this immediately leads to an expression for the sectional curvature of Σ_m^k. Le and Kendall [10] use this technique to give explicit expressions for the sectional and Ricci curvatures in terms of the singular values of the pre-shape matrix.

2.2. **Examples.**

2.2.1. $\mathbf{\Sigma_2^k}$. For k points in $\mathbb{R}^2$, we can exploit the algebra of the complex plane to simplify things. The Procrustes metric can be written in complex coordinates as

$$d(\pi(z), \pi(w)) = \inf_\theta(\cos^{-1}(\mathcal{R}e\{\sum_{j=1}^k z_j(e^{-i\theta}w_j^\star)\}))$$

where $w^\star$ is the complex conjugate. We can explicitly find the infimum as

$$d(\pi(z), \pi(w)) = \cos^{-1}(|\sum_{j=1}^k z_j w_j^\star|) = \cos^{-1}(\sqrt{\langle z, w\rangle^2})$$

In fact, since complex multiplication takes care of both rotations and scalings, we can identify Σ_2^k with $\mathbb{CP}^{k-2}$ in the following way. Let $z = (x_1, x_2, \ldots) \in \mathbb{C}^k$. If we remove translation $(x_1 - \bar{x}, x_2 - \bar{x}, \ldots)$ and then identify complex multiples $[z] = \{(\lambda(x_1 - \bar{x}), \lambda(x_2 - \bar{x}), \ldots)\ \forall\lambda\}$ we are left with exactly the set of complex lines through the origin in $\mathbb{C}^{k-1}$.

A limiting argument[7] starting from the Procrustes distance gives the line element

$$ds^2 = \frac{\langle z, z\rangle\langle dz, dz\rangle - \langle z, dz\rangle\langle dz, z\rangle}{\langle z, z\rangle^2}$$

which is, up to a constant factor, the standard *Fubini-Study* metric for $\mathbb{CP}^{k-2}$. The constant factor of course comes from the fact that our submersion endows $\mathbb{CP}^{k-2}$ with a constant sectional curvature of 4.

2.2.2. $\mathbf{\Sigma_2^3}$. For three points living in the complex plane, we can use a similarity transform to map two points to -1 and 1. The new position of the third point

$$z = \frac{2x_3 - (x_1 + x_2)}{x_2 - x_1}$$

encodes the shape of the triangle. To deal with the case that x_1 and x_2 lie on top of each other, we include the point at infinity which gives the appropriate limiting behavior for increasingly "skinny" triangles. This choice of standardization for triangles is refereed to as *Bookstein coordinates* after [2].

An easy argument shows that Σ_2^3 is isometric to the sphere of radius $\frac{1}{2}$. Stereographic projection of z maps from the closed complex plane to the sphere. Since the

Procrustes distance for Σ_2^n is the inverse cosine of a positive value, it is restricted to lie in $[0, \frac{\pi}{2}]$, hence we must project onto $\mathbb{S}^2(\frac{1}{2})$. This checks with our argument above, $\Sigma_2^3 \simeq \mathbb{CP}^1 \simeq \mathbb{S}^2(\frac{1}{2})$, as well as our curvature calculation for the Riemannian submersion. In this case our submersion $\pi : \mathbb{S}^3(1) \to \mathbb{S}^2(\frac{1}{2})$ takes us from a space of unit Gaussian curvature to one with a constant Gaussian curvature of 4.

2.2.3. $\mathbf{\Sigma_1^3}$. The set of degenerate triangles Σ_1^3 is simply the sphere $\mathbb{S}^1$. This sphere is immersed in the Bookstein model of Σ_2^3 as the image of the real axis under stereographic projection. How did we get from a circle of radius 1 to a radius of $\frac{1}{2}$? When we go from 1 to 2 dimensions, collinear triangles which were distinct as reflected copies of each other now become the same shape. Since reflected copies are antipodal in pre-shape space, identifying them takes $\Sigma_1^3 \simeq \mathbb{S}^1$ to $\mathbb{RP}^1$ which is then isometric to $\mathbb{S}^1(\frac{1}{2})$.

2.3. **Other groups acting on Σ_m^k.** It is worthwhile to consider what groups act on Σ_m^k. We begin by describing an action of $\mathbf{GL}(k-1)$ on Σ_m^k which commutes with rotation and scaling. Because centering removes a degree of freedom, it will be a convenient reminder of this fact to represent our pre-shape as a set of $k-1$ points. Let the *Helmert sub-matrix* $H \in \mathbb{R}^{k\times(k-1)}$ be defined as

$$h_{ij} = \begin{cases} -1/\sqrt{j(j+1)}, 1 \le i \le j \\ \sqrt{j/(j+1)}, i = j+1 \\ 0, i > j+1 \end{cases}$$

For a given point set X the *reduced point set*, defined as $\tilde{X} = XH \in \mathbb{R}^{k-1\times m}$, determines the original point set up to translation. To see this we note that the columns of H are a complete orthonormal set orthogonal to $(1, 1, \ldots, 1)$ so that $(X + \mathbf{w}\mathbf{1}_k^T)H = XH$. The action of $A \in \mathbf{GL}(k-1, \mathbb{R})$ on a reduced point set $\tilde{X}$ is given by right multiplication.

$$A \cdot \tilde{X} = \tilde{X}A$$

This clearly commutes with left-multiplication by rotations and translations of our original point set X.

Since the reduced point set is subsequently normalized for scale, the action of diagonal matrices in $\mathbf{GL}(k-1)$ descend to trivial actions on Σ_k^n. $\mathbf{O}(k-1)$ is a group of isometries with respect to the round metric on the pre-shape sphere. Our use of H above acts as a particular choice of isometry. The subgroup of permutations (corresponding to matrices with precisely a single 1 in each row and column) corresponds to re-labeling the points and is particularly useful in cases where the labels are incidental to some application. For example, Kendall uses this symmetry extensively (c.f. [7]) in order to give a compact visual display of probability densities on Σ_2^3.

When m is even, there is an additional symmetry, *conjugation*, which is represented as multiplication on the *left* by $diag(1, 1, \ldots, 1, -1)$. For m odd, conjugation is equivalent (modulo $\mathbf{SO}(m)$) to $-\mathbf{I} \in \mathbf{GL}(k-1)$.

Are there other linear transformations $T : \mathbb{R}^{n\times(k-1)} \to \mathbb{R}^{n\times(k-1)}$ acting on reduced point sets which descend to shape space? Swann and Olsen [16] answer in the negative

Theorem 2.3 (Swann and Olsen, 2003)**.** *For $m > 2$ the maximal connected group of linear transformations on reduced point sets which induce an action on shape space is $\mathbf{SO}(m) \times \mathbf{GL}(k-1)$. For $m = 2$ the group is $\mathbf{GL}(k-1, \mathbb{C})$.*

2.4. **Topology of Σ_m^k.** We have already seen that several of the low-dimensional shape spaces are homeomorphic to spheres. It is also the case that the *simplicial* shape spaces Σ_m^{m+1} are homeomorphic to spheres. Note that the submersion bounds on the curvature aren't strong enough to invoke the $\frac{1}{4}$-pinching topological sphere result (see e.g. [1]). In fact, the shape space $\Sigma_2^4 \simeq \mathbb{CP}^2$ provides the usual example showing that that strict $\frac{1}{4}$-pinching is optimal.

The first proof that the simplicial shape spaces are homeomorphic to spheres is attributed to unpublished work by A. J. Casson (c.f. [7]). A different proof by Le[9] begins with an explicit description of the cut locus for Σ_k^m when $2 \leq m \leq k-1$ which we will now summarize.

Consider the distance between shapes $\pi(X)$ and $\pi(Y)$

$$d(\pi(X), \pi(Y)) = \inf_{T \in \mathbf{SO}(m)} d(X, TY) = \cos^{-1}\left(\sup_{T \in \mathbf{SO}(m)} Tr(TYX^T)\right)$$

The singular value decomposition of $YX^T = U\Lambda V^T$ with $\Lambda = diag\{\lambda_1 \geq \ldots \lambda_{m-1} \geq |\lambda_m|\}$ and $U, V \in \mathbf{SO}(m)$ gives the distance as

$$d(\pi(X), \pi(Y)) = \cos^{-1}\left(\sup_{T \in \mathbf{SO}(m)} Tr(V^T T U \Lambda)\right) = \cos^{-1}\left(\sum \lambda_i\right)$$

which is achieved when $T^T = UV^T$. The point $U^T V Y$ is a point on the fiber above $\pi(Y)$ whose distance to X is as small as possible. However, the optimal rotation, and hence the closest point, can fail to be unique if $\lambda_m + \lambda_{m-1} \leq 0$ (i.e. $rank(YX^T) < m-1$ or $\lambda_m + \lambda_{m-1} = 0$). Le shows that these two cases completely describe the cut locus for a non-singular point.

Theorem 2.4 (Le, 1991)**.** *The cut locus on Σ_m^k for any non-singular point $\pi(X)$ is $\{\pi(Y) : rank(YX^T) < m-1\} \cup \{\pi(Y) : \lambda_m = \lambda_{m-1} < 0\}$*

From this point the exponential map centered at two poles of the pre-shape sphere composed with π takes the closed balls in the tangent space homeomorphically to closed balls that cover Σ_m^{m+1}. Identifying corresponding points on the two boundaries we arrive at

Theorem 2.5 (Le, 1991)**.** Σ_m^{m+1} *($m \geq 3$) is homeomorphic but not isometric to* $S^{\frac{1}{2}m(m+1)-1}$

This result provides further insight into the *over-dimensioned* spaces $\{\Sigma_m^k : k < m+1\}$. In this case, an appropriate element of $\mathbf{SO}(m)$ will render the bottom $m-(k-1)$ rows of a pre-shape X equal to 0 hence identifying pre-shapes of such an Σ_m^k with Σ_{k-1}^k. Of course distinct conjugate shapes in the simplicial shape space Σ_{k-1}^k that lie in complementary hemispheres are identified in the over-dimensioned space Σ_m^k by a rotation through the extra dimensions. Since the simplicial shape space is homeomorphic to a sphere, the over-dimensioned space is homeomorphic to a hemisphere and hence contractible.

3. The shapes of planar curves

We will now switch gears and look at something closer to Riemann's infinite-dimensional "manifoldness of possible shapes". The results here are somewhat less developed due to the added difficulties of treating infinite-dimensional manifolds. We begin with a function space describing plane curves and then move to a slightly more complicated space implicitly described by pairwise relationships between curves. Lastly, we point out a third approach for measuring distance between curves using diffeomorphisms of the plane.

3.1. **The tangent space to direction functions.** A closed planar curve parameterized by arc length can be represented most directly as $\gamma : [0, 2\pi] \to \mathbb{R}^2$. The direction function θ is set by $\dot{\gamma}(s) = \exp(i\theta(s))$ and curvature defined by $\kappa(s) = \dot{\theta}(s)$. We will initially consider representing shapes using the direction function θ. Take θ to be an absolutely continuous, square integrable function defined a.e. To assure invariance under rotation, we subtract a constant so that θ has mean value π:

$$\phi_1 = \frac{1}{2\pi}\int_0^{2\pi} \theta(s)ds = \pi$$

This choice of normalization gives $\theta(s) = s$ for the unit circle. To assure that θ describes a closed curve, it must satisfy

$$\phi_2 = \int_0^{2\pi} \cos(\theta(s))ds = 0$$
$$\phi_3 = \int_0^{2\pi} \sin(\theta(s))ds = 0$$

These constraints are easily seen to pick out a sub-manifold of L^2 using the map $\phi = (\phi_1, \phi_2, \phi_3) : L^2 \to \mathbb{R}^3$. To see that that $d\phi$ is surjective write

$$d\phi_1 f = \frac{1}{2\pi}\int_0^{2\pi} f ds = \langle f, \frac{1}{2\pi}\rangle$$
$$d\phi_2 f = -\int_0^{2\pi} \sin(\theta) f ds = -\langle f, \sin(\theta)\rangle$$
$$d\phi_3 f = \int_0^{2\pi} \cos(\theta) f ds = \langle f, \cos(\theta)\rangle$$

We call the pre-image $\mathcal{P}_C = \phi^{-1}(\pi, 0, 0)$ the *pre-shape space* because multiple elements of $\mathcal{P}_C$ have the same shape but are parameterized from a different starting point $s = 0$. We give the action of $g \in \mathbf{SO}(2)$ as $(g\theta)(s) = \theta((s - g)\, mod\, 2\pi)$ and take as our shape space $\mathcal{S}_C = \mathcal{P}_C/\mathbf{SO}(2)$. A vector $f \in L^2$ is tangent to $\mathcal{P}_C$ at θ only if it is orthogonal to the normal space spanned by $\{1, \sin(\theta), \cos(\theta)\}$ so the tangent space is given by

$$T_\theta(\mathcal{P}_C) = \{f \in L^2 | f \perp span\{1, \sin(\theta), \cos(\theta)\}\}$$

Define a metric on $\mathcal{S}_C$ by taking the minimum distance in $\mathcal{P}_C$ between the orbits of two points under the re-parameterizations. Since $\mathbf{SO}(2)$ acts on $\mathcal{P}_C$ by isometries, we can fix one point and look for nearest point in the orbit of the other. Such a geodesic from point p_1 must be orthogonal to the $\mathbf{SO}(2)$ orbit of p_2. In fact, it is orthogonal to every $\mathbf{SO}(2)$ orbit it meets. Since $T_\theta(S^1(\theta))$ is spanned by $1 - \dot{\theta}$ the tangents to the geodesic must be orthogonal to $1 - \dot{\theta}$. The authors of [8] suggest algorithms for finding geodesics in this space by first estimating an initial tangent vector using a Fourier approximation and then alternately taking a small step in the tangent direction and projecting the resulting function from L^2 into $\mathcal{S}_C$.

3.2. **Correspondence.** One difficulty that plagues our understanding of continuous shapes is that of correspondence. In the case of point sets, each point was given a label and two shapes were compared by using the labels to map them into the pre-shape sphere in a canonical way (permutations of the labels yield alternate isometries). For non-rigid shapes like plane curves, the notion of corresponding "parts" is far less precise. For example, the figure at right shows two curves with arrows marking points which intuitively "correspond". If we were asked to match the two curves, surely it would make sense that the pair of protrusions should be roughly lined up via a rigid motion and then the rest of the figure smoothly deformed as necessary.

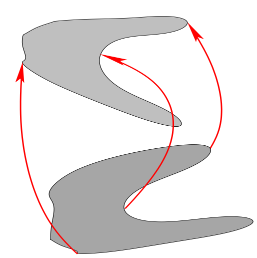

In an attempt to gain some mathematical precision, we follow [17] and define

Definition 3.1. A *correspondence* between any two sets U and V is a subset Φ of $U \times V$ such that for every $u \in U$ (resp. $v \in V$) there exists some $v \in V$ (resp. $u \in U$) such that $(u, v) \in \Phi$.

Of course, what we would really like is that Φ be a graph of some function (either from U onto V or V onto U) in every open neighborhood. In the case of two closed plane curves, this assures the desired topological picture that Φ be a connected, smooth, compact manifold embedded in the torus.

Definition 3.2. A correspondence between topological spaces U and V is *graph-like* if there is a family $\{(U_\alpha, V_\alpha, \Phi_\alpha) : \alpha \in \mathcal{A}\}$ such that $\{U_\alpha : \alpha \in \mathcal{A}\}$ is an open cover of X, $\{V_\alpha : \alpha \in \mathcal{A}\}$ is an open cover of Y, $\Phi = \cup_\alpha \Phi_\alpha$, and for each $\alpha \in \mathcal{A}$, either Φ_α is the graph of an onto function from U_α to V_α or from V_α to U_α

Definition 3.3. A *bimorphism* between simple regular closed curves C and D is a graph-like, differentiable correspondence in which the image of every point in C or D is connected.

A bimorphism Φ is itself a regular curve with some parameterization $p : (0, 1] \to \Phi$ and is covered by two domains $\Phi_C = \{p(t) : (\pi_C \circ p(t))' \neq 0\}$ and $\Phi_D = \{p(t) : (\pi_D \circ p(t))' \neq 0\}$ on which the respective projections π_C and π_D are diffeomorphisms. We define an arc length for bimorphisms as the weighted sum of the Euclidean length of projections

$$ds = \sqrt{\left\| \frac{(\pi_C \circ p)'}{l_C} \right\|^2 + \left\| \frac{(\pi_D \circ p)'}{l_D} \right\|^2} dt$$

where l_C, l_D are the lengths of C and D. Now let $\gamma(t) : (0, l_\Phi] \to \Phi$ be the arc length parameterization of Φ. To measure the difference in local shape between corresponded points $\pi_C(\gamma(t))$ and $\pi_D(\gamma(t))$ we compare the derivative of the direction function at the points

$$G(t) = \theta_C(\pi_C(\gamma(t)))' - \theta_D(\pi_D(\gamma(t)))'$$

For some choice of non-negative, symmetric function F we may define the distance between curves as

$$d(C, D) = \min_{\Phi} \int_0^{l_\Phi} F(G(t; \Phi))dt$$

Note that weighting by lengths in the definition of ds makes this distance invariant under scaling of C or D while the minimization over Φ gives invariance under rotation.

By integrating G, [17] show that this distance connects up in a nice way with the Gauss map. In particular, the singularities of the Gauss map break suitable curves into a collection of segments. This operation could provide an interesting way to operationalize the intuition of "parts" described above.

3.3. **Groups of Diffeomorphisms.** A third way of looking at curves it to consider diffeomorphisms of the entire plane which carry along the curve of interest. This approach has the advantage of extending to other sorts of objects beyond closed curves and was originally pursued by Grenander for solving the problem of matching images to templates (see e.g. [6]). We will begin with a general procedure for passing between a deformation cost on our space of shapes and a distance on a diffeomorphism group.

Given a group $\mathbf{G}$ which acts transitively on some space of shapes $\mathcal{S}$, let the deformation cost $\Gamma : \mathbf{G} \times \mathcal{S} \to [0, \infty)$ satisfy $\Gamma(e, C) = 0$, $\Gamma(a, C) = \Gamma(a^{-1}, a \cdot C)$ and $\Gamma(ab, C) \leq \Gamma(b, C) + \Gamma(a, bC)$ for $a, b \in \mathbf{G}, C \in \mathcal{S}$. Then the function d

$$d(C_1, C_2) = \inf\{\Gamma(a, C_1), C_2 = a \cdot C_1\}$$

is symmetric, satisfies the triangle inequality and $d(C, C) = 0$. When $\mathbf{G}$ acts transitively on $\mathcal{S}$ we can identify $\mathcal{S}$ with the quotient of $\mathbf{G}$ by the isotropy group of some reference element C_0. This gives us the following proposition[19]

Proposition 3.4. Let d_G be a symmetric distance on the group $\mathbf{G}$ so that d_G satisfies the triangle inequality and $d_G(g, g) = 0$. Assume there exists $\gamma_0 : \mathbf{G} \to \mathbb{R}$ such that $\gamma_0(h) = 1$ whenever $h \cdot C_0 = C_0$ and for all $f, g, h \in \mathbf{G}$ we have $d_G(hf, hg) = \gamma_0(h) d_G(f, g)$. Then for any two objects $C_1 = f \cdot C_0$, $C_2 = g \cdot C_0$ the deformation cost

$$\Gamma_0(h, C_1) = \frac{1}{\gamma_0(f)} d_G(e, h^{-1})$$

induces the distance

$$d_0(C_1, C_2) = \frac{1}{\gamma_0(f)} \inf_h \{d_G(e, h), h \cdot C_2 = C_1\}$$

on $\mathcal{S}$. A different choice, C_0', of reference object yields a cost function which differs by a multiplicative factor.

Using this technique, Younes[19] gives the following construction for (possibly open) plane curves. The idea is to start with an intuitive variational cost and compute an expression for $d_G(e, a)$ when a is near the identity e. Younes then works backward by defining a group action on a subset of Hilbert space rendering it homomorphic to G and showing that the usual L^2 norm agrees with the desired d_G.

Consider an infinitesimal deformation of a plane curve $C = \{(x(t), y(t)) : t \in [0, l]\}$ parameterized by arc-length. Such a deformation is described by a vector field

V which gives the infinitesimal displacement $C \mapsto V \cdot C \equiv (x(t)+u(t), y(t)+v(t))$. We define the energy (squared cost) of this deformation as

$$\delta E(V) = \int_0^l \|\dot{V}(t)\|^2 dt$$

Note that the energy for a translation of the plane is 0. Let l_V be the length of our curve after the deformation. It is convenient to separate the action of V on C into an arc-length term $\tilde{g}_V : [0, l] \to [0, l_V]$ which gives the arc length of $V \cdot C$ for each point on C and a direction term $\tilde{\theta}_V : [0, l] \to [0, 2\pi)$. To order 1, we have the tangential component

$$T = \dot{\tilde{g}}_V - 1 = \dot{u}(t)\dot{x}(t) + \dot{v}(t)\dot{y}(t)$$

and the normal component

$$N = \tilde{\theta}_V \circ g_V - \theta = -\dot{y}(t)\dot{u}(t) + \dot{x}(t)\dot{v}(t)$$

so we can write

$$\delta E(V) = \int_0^l (\dot{\tilde{g}}_V - 1)^2 dt + \int_0^l (\tilde{\theta}_V \circ \tilde{g}_V - \tilde{\theta})^2 dt$$

which is simplified by substituting $g_V(t) = (1/l)\tilde{g}_V(lt)$, $\theta_V(t) = \tilde{\theta}_V(lt)$ and $\lambda = l_V/l$ and writing

$$\delta E(\lambda, g_V, N; l, \theta) = l \int_0^1 (\lambda \dot{g}_V - 1)^2 + N^2 dt$$

where (λ, g_V, N) transforms a curve specified by (l, θ) to a new curve $(l/\lambda, \theta \circ g_V^{-1} + N \circ g_V^{-1})$.

To tighten our notation, let Ω be the unit circle in $\mathbb{C}^1$. Let our class of curves be described (modulo translations) by

$$\mathcal{S} = \{(l, \eta) : l > 0, \eta : [0, 1] \to \Omega\}$$

Our group

$$\mathbf{G} = \{(\lambda, g, r) : \lambda > 0, g : [0, 1] \to [0, 1], r : [0, 1] \to \Omega\}$$

where g is a diffeomorphism, acts on $\mathcal{S}$ by

$$(\lambda, g, r) \cdot (l, \eta) = (l/\lambda, r\eta \circ g)$$

and group multiplication is set by

$$(\lambda_1, g_1, r_1)(\lambda_2, g_2, r_2) = (\lambda_2\lambda_1, g_2 \circ g_1, r_1(r_2 \circ g_1))$$

with inverse $(1/\lambda, g^{-1}, r^\star \circ g^{-1})$. We will let our squared deformation cost be the the first order terms of δE, so that for a curve (l, η):

$$\Gamma((\lambda, g, r)^{-1}, (l, \eta))^2 = l(\lambda - 1)^2 + \int_0^1 (\dot{g}(t) - 1)^2 + |r(t) - 1|^2 dt$$

If we take our reference curve to be the horizontal segment $C_0 = (1, \mathbf{1})$, then by Proposition 3.4 when $C = (l, \eta) = (\lambda, g, r) \cdot C_0$ we have that $\lambda = 1/l$ and $r = \eta$ so $\gamma_0((\lambda, g, r)) = 1/\sqrt{l}$ and for a group element (λ, g, r) which is infinitesimally close the the identity, the group distance is given by

$$d_G(e, (\lambda, g, r))^2 = (\lambda - 1)^2 + \int_0^1 (\dot{g}(t) - 1)^2 + |r(t) - 1|^2 dt$$

To extend this distance to all of G, Younes constructs a group homomorphism from $\tilde{\mathbf{G}} = \{X \in L^2([0,1],\mathbb{C}) : |X| > 0 \quad a.e.\}$ to $\mathbf{G}$ with the product on $\tilde{\mathbf{G}}$ given by

$$(X \cdot Y)(s) = X(s)Y(X(s)/\|X\|_2^2)$$

The standard L^2 norm on $\tilde{\mathbf{G}}$ agrees with the infinitesimal d_G derived above and results in the following concrete distance between plane curves $C_1 = (l_1, e^{i\theta_1})$ and $C_2 = (l_2, e^{i\theta_2})$

$$d(C_1,C_2)^2 = l_1 + l_2 - 2\sqrt{l_1 l_2}\sup_g \int_0^1 \sqrt{\dot{g}(t)}\left|\cos\frac{\theta_2 \circ g - \theta_1}{2}\right| dt$$

The supremum is over all increasing diffeomorphisms of $[0,1]$ (in analogy to our earlier minimization over bimorphisms). As noted[19], $\mathbf{G}$ does not act on the subset of *closed* plane curves $\mathcal{S}_C \subset \mathcal{S}$ (except for trivial $g \in \mathbf{G}$). Unlike the geodesic paths described in Section 3.1, there is nothing to stop the evolving curve from "breaking" along the geodesic deformation from closed curve C_1 to closed curve C_2.

For further reading on deformation based metrics, the reader is referred to [18] where the authors consider deforming both the range and domain of some smooth function. [5] study the well posedness of a variational problem which trades off the amount of deformation with a direct (straight-line) matching distance.

Appendix: Riemannian Submersions

Theorem 3.5 (O'Neill, 1966). *If M is submersed into $\tilde{M}$, then the curvatures are related by:*

$$\langle \tilde{R}(\tilde{X},\tilde{Y})\tilde{Z},\tilde{W}\rangle = \langle R(X,Y),Z,W\rangle - \frac{1}{4}\langle[\tilde{X},\tilde{Z}]^v,[\tilde{Y},\tilde{W}]^v\rangle - \frac{1}{4}\langle[\tilde{Y},\tilde{Z}]^v,[\tilde{X},\tilde{W}]^v\rangle - \frac{1}{2}\langle[\tilde{X},\tilde{Y}]^v,[\tilde{Z},\tilde{W}]^v\rangle$$

Proof. We proceed in 4 steps

P0: Let V be vertical, then using the expression $2\langle\nabla_X Y, Z\rangle = X\langle Y,Z\rangle + Y\langle Z,X\rangle - Z\langle X,Y\rangle + \langle[X,Y],Z\rangle - \langle[X,Z],Y\rangle - \langle[Y,Z],X\rangle$ we have that

$$\begin{aligned} 2\langle\tilde{\nabla}_{\tilde{X}}V,\tilde{Y}\rangle &= Y\langle\tilde{Z},\tilde{X}\rangle - \langle[\tilde{X},\tilde{Z}],V\rangle \\ &= -(-Y\langle\tilde{Z},\tilde{X}\rangle + \langle[\tilde{X},\tilde{Z}],V\rangle) \\ &= -2\langle\tilde{\nabla}_{\tilde{X}}\tilde{Y},V\rangle \end{aligned}$$

P1: Let V be vertical, then $[\tilde{X},V]$ is also vertical. Using again the expression for $\langle\nabla_X Y,Z\rangle$ and the orthogonality of horizontal and vertical vector fields, we have that $2\langle\nabla_X Y,V\rangle = -Z\langle X,Y\rangle + \langle[X,Y],Z\rangle = \langle[X,Y],Z\rangle$ since $\langle X,Y\rangle$ is constant in the vertical direction. For an arbitrary vector field, this gives the decomposition

$$\begin{aligned} \langle\tilde{\nabla}_{\tilde{X}}\tilde{Y},\tilde{H}+V\rangle &= \langle\tilde{\nabla}_{\tilde{X}}\tilde{Y},\tilde{H}\rangle + \langle\tilde{\nabla}_{\tilde{X}}\tilde{Y},V\rangle \\ &= \langle\widetilde{\nabla_X Y,H}\rangle + \frac{1}{2}\langle[\tilde{X},\tilde{Y}],V\rangle \end{aligned}$$

P2: Using **P1** twice and properties of ∇

$$\begin{aligned}
\langle \tilde{\nabla}_{\tilde{X}} \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{W} \rangle &= \tilde{X} \langle \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{W} \rangle - \langle \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{\nabla}_{\tilde{X}} \tilde{W} \rangle \\
&= \tilde{X} \langle \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{W} \rangle - \langle \widetilde{\nabla_Y Z} + \frac{1}{2}[\tilde{Y}, \tilde{Z}]^v, \widetilde{\nabla_X W} + \frac{1}{2}[\tilde{X}, \tilde{W}]^v \rangle \\
&= \tilde{X} \langle \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{W} \rangle - \langle \widetilde{\nabla_Y Z}, \widetilde{\nabla_X W} \rangle - \frac{1}{4} \langle [\tilde{Y}, \tilde{Z}]^v, [\tilde{X}, \tilde{W}]^v \rangle \\
&= \langle \nabla_X \nabla_Y Z, W \rangle - \frac{1}{4} \langle [\tilde{Y}, \tilde{Z}]^v, [\tilde{X}, \tilde{W}]^v \rangle
\end{aligned}$$

P3: Using **P0**, $\langle \tilde{\nabla}_V \tilde{X}, \tilde{Y} \rangle = \langle [V, \tilde{X}], \tilde{Y} \rangle + \langle \tilde{\nabla}_{\tilde{X}} V, \tilde{Y} \rangle = -\langle \tilde{\nabla}_{\tilde{X}} \tilde{Y}, V \rangle$. Now split $[\tilde{X}, \tilde{Y}]$ into vertical and horizontal parts and to get

$$\begin{aligned}
\langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]} \tilde{Z}, \tilde{W} \rangle &= \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]^h} \tilde{Z}, \tilde{W} \rangle + \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]^v} \tilde{Z}, \tilde{W} \rangle \\
&= \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]^h} \tilde{Z}, \tilde{W} \rangle - \langle [\tilde{X}, \tilde{Y}]^v, \tilde{\nabla}_{\tilde{Z}} \tilde{W} \rangle \\
&= \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]^h} \tilde{Z}, \tilde{W} \rangle - \langle [\tilde{X}, \tilde{Y}]^v, \widetilde{\nabla_Z W} + \frac{1}{2}[\tilde{Z}, \tilde{W}] \rangle \\
&= \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]^h} \tilde{Z}, \tilde{W} \rangle - \frac{1}{2} \langle [\tilde{X}, \tilde{Y}]^v, [\tilde{Z}, \tilde{W}]^v \rangle
\end{aligned}$$

Combining **P2** and **P3** we can write down the curvature

$$\begin{aligned}
&\langle \tilde{R}(\tilde{X}, \tilde{Y}) \tilde{Z}, \tilde{W} \rangle \\
&= \langle \tilde{\nabla}_{\tilde{Y}} \tilde{\nabla}_{\tilde{X}} \tilde{Z}, \tilde{W} \rangle - \langle \tilde{\nabla}_{\tilde{X}} \tilde{\nabla}_{\tilde{Y}} \tilde{Z}, \tilde{W} \rangle + \langle \tilde{\nabla}_{[\tilde{X}, \tilde{Y}]} \tilde{Z}, \tilde{W} \rangle \\
&= \langle R(X, Y), Z, W \rangle - \frac{1}{4} \langle [\tilde{X}, \tilde{Z}]^v, [\tilde{Y}, \tilde{W}]^v \rangle - \frac{1}{4} \langle [\tilde{Y}, \tilde{Z}]^v, [\tilde{X}, \tilde{W}]^v \rangle \\
&\quad - \frac{1}{2} \langle [\tilde{X}, \tilde{Y}]^v, [\tilde{Z}, \tilde{W}]^v \rangle
\end{aligned}$$

□

References

[1] U. Abresch, W. Meyer "Injectivity Radius Estimates and Sphere Theorems" *Comparison Geometry* MSRI Publications, Vol. 30, 1997

[2] F. L. Bookstein "Size and shape spaces for landmark data in two dimensions" *Statistical Science* Vol. 1, pp 181–242, 1986

[3] T. K. Carne "The geometry of shape spaces" *Proceedings of the London Mathematical Society* Vol. 61, pp 407–432, 1990

[4] I. Dryden, K. Mardia *Statistical Shape Analysis.* John Wiley & Sons, 1998

[5] P. Dupuis, U. Grenander, M. Miller "A Variational Formulation of a Problem in Image Matching" *Quarterly of Applied Mathematics* Vol. 56, pp. 587-600, 1998

[6] U. Grenander *Lectures in Pattern theory* Springer Verlag, Berklin, 1976

[7] D. Kendall "Shape manifolds, procrustean metrics, and complex projective spaces," *Bull. London Math. Soc.* Vol. 16, pp. 81–121, 1984.

[8] E. Klassen, A. Srivastava, W. Mio "Analysis of Planar Shapes Using Geodesic Paths on Shape Spaces" *IEEE Transactions on Pattern Analysis and Machine Intelligence* In press.

[9] H. Le "On geodesics in Euclidean shape spaces," *J. London Math. Soc.* Vol. 44, pp. 360–372, 1991.

[10] H. Le, D. Kendall, "The Riemannian Structure of Euclidean Shape Spaces: A Novel Environment for Statistics," *The Annals of Statistics*, Vol. 21, No. 3, pp. 1225–1271, 1993.

[11] M. Miller, L. Younes "Group Actions, Homeomorphisms, and Matching: A general framework" *International Journal of Computer Vision* Vol. 41, pp 61–84, 2001
[12] R. Montgomery "The geometric phase of the three-body problem" *Nonlinearity* Vol. 9, pp 1341-1360, 1996.
[13] B. O'Neill "The fundamental equations of a submersion" *Michigan Mathematics Journal* Vol. 13, pp 459–469, 1966
[14] B. Riemann (trans. W. Clifford) "On the Hypotheses which lie at the Bases of Geometry," *Nature* Vol. 8, Nos. 183,184, pp. 14–17,36,37
[15] C. Small *The Statistical Theory of Shape* Springer, 1996.
[16] A. Swann, N. Olsen "Linear Transformation Groups and Shape Space" *Journal of Mathematical Imaging and Vision* Vol. 19, pp 49–62, 2003.
[17] H. Tagare, D. O'Shea, D. Groisser "Non-Rigid Shape Comparison of Plane Curves in Images" *Journal of Mathematical Imaging and Vision* Vol. 16, pp 57–68, 2002
[18] A. Trouvé, L. Younes "Local analysis on a shape manifold" Preprint no. 2002-03 du laboratoire LAGA, Universite Paris 13, 2002
[19] L. Younes "Computable elastic distance between shapes" *SIAM Journal on Applied Mathematics* Vol. 58, No 2, pp 565–586, 1997

Current address: Department of Electrical Engineering and Computer Science University of California, Berkeley, CA 94720

E-mail address: `fowlkes@cs.berkeley.edu`