

Locally-Weighted Elastic Comparison of Planar Shapes

Justin Strait, Sebastian Kurtek, and Steven MacEachern
Department of Statistics
The Ohio State University

strait.50@osu.edu, kurtek.1@stat.osu.edu, snm@stat.osu.edu

Abstract

Registration of curves is a necessary component of statistical shape analysis. The goal of registration is to align collections of shapes so that common features are appropriately matched for further comparison and subsequent analyses. Traditional methods for registration typically rely on optimizing an energy functional over a set of appropriate shape-preserving transformations (i.e., rotations and re-parameterizations). These functionals typically rely on the standard \mathbb{L}^2 metric. In certain applications, it may make sense to use a more flexible metric which can align shapes most preferably with respect to a local shape feature (i.e., a certain curve segment selected from the overall shape). In this work, we define a weighted shape metric which allows for emphasis on local shape features. Registration can be performed with respect to this metric. We demonstrate the registration procedure using simulated curves as well as real data, and show the dependence of the optimal rotation and re-parameterization on the specified weights, as well as the resulting deformation path from one shape to another.

1. Introduction

Statistical shape analysis is the study of data sets which consist of outlines of objects, generally extracted from images. This type of data is available in numerous applications, including computer vision, biology, anatomy, medical imaging, biometrics, and forensic analysis. A major challenge in shape analysis is establishing equivalence of shape, which can be defined as a property which is unchanged after applying various shape-preserving transformations (generally translation, scale, and rotation). In order to analyze shape data, one must respect these invariances, as objects which appear different can still be identically-shaped. Thus, developing a suitable shape representation is crucial to shape analysis. Once found, a metric between shapes is defined, allowing for further statistical inference.

The two most general classes of shape representations are landmark-based and function-based. Landmark meth-

ods represent a shape using a finite set of labeled points, known as landmarks [6, 4, 12, 2]. These points are selected mathematically (e.g., points of extreme curvature) and/or anatomically (as meaningful points to the function of the object of interest). This allows for a finite-dimensional approximation of shape, and opens up multivariate analysis techniques for statistical inference. More recently, researchers have worked on function-based methods, since an object's shape is generally thought of as a function rather than a discrete set of points. By treating the outline of an object as a continuous function, shape can also be thought of as being preserved even if the function is re-parameterized. Parameterization-invariance can be introduced by standardizing all parameterizations to arc-length; however, this has been shown to be sub-optimal, as it forces a correspondence of features between shapes that may not necessarily be appropriate [20, 7]. An alternative approach is to optimize over the space of valid re-parameterization functions. This allows for more appropriate shape matching (known as registration), and underlies what are known as elastic representations of shape [19, 10, 13].

Recently, Strait *et al.* [18] combined the two approaches discussed into a landmark-constrained elastic shape representation (also see [1, 9] for related methods). This allows for treatment of an object's outline as a function, while also providing for subject expertise input in the way of annotated landmarks on the outline. In medical imaging, landmarks often represent important anatomical features to the structure; marking them places extra emphasis on these features of interest. Strait *et al.* [18] showed improvements in statistical performance in cases where landmarks provide useful information. Marking landmarks on an outline segments the shape into "pieces", which are then used for subsequent registration steps. Registration relies on the \mathbb{L}^2 metric between square root velocity functions (defined in [13]), which is rewritten into a sum of integrals over each shape segment. Naturally, one might wonder if allowing for flexible weighting of each of these integrals over particular segments can further improve statistical performance. By up-weighting the outline of a bird's beak, perhaps one can see

improvements in identifying bird species with vastly different beak structures. Our goal is to formulate a method for registration of full shapes dominated by smaller-scale (i.e., local) features, and induce a distance which scrutinizes the highly-weighted local features of interest.

In this paper, we introduce a metric which allows for flexible weighting of local features, based on the landmark-constrained representation in [18]. In Section 2, we outline the landmark-constrained representation and introduce the weight function. We show how this impacts registration of shapes as well as the shape metric and geodesic. Section 3 demonstrates weighted registration on four sets of simulated curves. We then apply this new framework to three sources of real shape data in Section 4, and conclude with summarizing remarks and future directions in Section 5.

2. Mathematical Framework

In this section, we describe the main shape representation framework used throughout this work; further details can be found in [13, 18]. Once the representation is developed, methods for shape registration and computation of a shape metric and geodesic are introduced.

2.1. Square Root Velocity Function

Let $\beta : \mathcal{D} \rightarrow \mathbb{R}^d$ be the absolutely continuous curve representing the object of interest's outline. \mathcal{D} is the curve domain, taken to be $[0, 1]$ for open curves and \mathbb{S}^1 for closed curves. For the time being, we will focus on open curves in $d = 2$ dimensions (i.e., planar curves); however, there are natural extensions of this method to closed curves and curves of higher dimension. A popular shape representation for β is the *square root velocity function (SRVF)* [13]:

$$q(t) = \begin{cases} \frac{\dot{\beta}(t)}{\sqrt{|\dot{\beta}(t)|}} & \text{if } \beta \text{ is differentiable at } t \text{ and } |\dot{\beta}(t)| \neq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where $|\cdot|$ is the Euclidean norm and $\dot{\beta}$ is the time-derivative of β . The SRVF has numerous benefits. First, q encodes all local information about β : the instantaneous speed and direction of β can be written explicitly in terms of q . In addition, given starting point $\beta(0)$ of β , there is a smooth bijection between q and β : $\beta(t) = \beta(0) + \int_0^t q(s)|q(s)| ds$. Since the SRVF is a function of $\dot{\beta}$, it is automatically invariant to translation. Scale invariance can be introduced by re-scaling β to unit length. In particular, if the unit length constraint $\int_0^1 |\dot{\beta}(t)| dt = \int_0^1 |q(t)|^2 dt = 1$ is imposed, the space of corresponding SRVFs is the infinite-dimensional unit Hilbert sphere, called the *pre-shape space* $\mathcal{C} = \{q : [0, 1] \rightarrow \mathbb{R}^2 \mid \int_0^1 |q(t)|^2 dt = 1\}$. If one wishes to analyze size-and-shape of an object, the re-scaling step is skipped, and SRVFs lie in the *pre-size-and-shape space*

$\mathcal{C}_s = \{q : [0, 1] \rightarrow \mathbb{R}^2\}$ (i.e., the space of all square integrable functions). For this work, we consider SRVFs in the ambient space \mathcal{C}_s . In order to proceed, we must introduce rotation and re-parameterization invariance.

2.2. Weighted Metric on \mathcal{C}_s

Suppose that we are provided an open curve β with k landmark locations marked, denoted $\{\beta(\theta_1), \dots, \beta(\theta_k)\} \in \mathbb{R}^2$. Without loss of generality, assume $0 < \theta_1 < \dots < \theta_k < 1$. Specifying k landmarks splits β into $k + 1$ curve segments. These can be identified by partitioning the curve domain \mathcal{D} : define $S_1 = [0, \theta_1), S_2 = [\theta_1, \theta_2), \dots, S_k = [\theta_{k-1}, \theta_k), S_{k+1} = [\theta_k, 1]$. Notice that these sets do form a partition, as $\mathcal{D} = \bigcup_{i=1}^{k+1} S_i$ and $S_i \cap S_j = \emptyset$ for $i \neq j$. Then, the i^{th} curve segment of β is the absolutely continuous function $\beta^{(i)} : S_i \rightarrow \mathbb{R}^2$ defined as the restriction $\beta^{(i)} := \beta|_{S_i}$ (with corresponding SRVF segment $q^{(i)}$). In practice, landmarks are specified manually by the researcher to capture important local shape features, using a discretized curve which is arc-length parameterized; see [18] for further details. In the case where landmark locations are unknown, automatic detection may be necessary; see [11, 16, 15, 3].

To weight curve segments, define weight function $w : [0, 1] \rightarrow \mathbb{R}$ by $w(t) = w_i \mathbb{1}_{\{t \in S_i\}}$, with $w_i > 0 \forall i$, $\sum_i w_i = k + 1$. This assigns weight w_i to segment $\beta^{(i)}$. Note that this weight function is discretely defined; a continuous weight function may be desirable in certain applications, and is left as future work. Also note the weight sum restriction, which is one way to standardize weights (and allows for equally weighted segments to be a special case). We require weights to be non-negative for mathematical reasons described below. Weights can be selected to emphasize local shape features that are of interest, and impacts shape registration, as described in the next section.

Given two curves β_1, β_2 with corresponding SRVFs $q_1, q_2 \in \mathcal{C}_s$ and *common* landmark locations $\theta_1, \dots, \theta_k$, define the weighted \mathbb{L}^2 inner product between their SRVFs,

$$\begin{aligned} \langle\langle q_1, q_2 \rangle\rangle^{(w)} &:= \int_0^1 w(t) \langle q_1(t), q_2(t) \rangle dt \\ &= \sum_{i=1}^{k+1} w_i \int_{S_i} \langle q_1(t), q_2(t) \rangle dt, \end{aligned} \quad (2)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. This induces the *weighted landmark-constrained pre-size-and-shape (WLCP) metric*:

$$d_{\mathcal{C}_s}^{(w)}(q_1, q_2) = \left(\sum_{i=1}^{k+1} w_i \int_{S_i} |q_1(t) - q_2(t)|^2 dt \right)^{1/2}, \quad (3)$$

Notice that when $w_i = 1 \forall i$, $d_{\mathcal{C}_s}^{(w)}(q_1, q_2) = d_{\mathcal{C}_s}(q_1, q_2)$ (i.e., the original, unweighted \mathbb{L}^2 metric defined in [18]).

The WLCP metric satisfies required properties for distance functions only if w_i are strictly non-negative for all i .

A question which arises is why this metric involves SRVFs q_1, q_2 , rather than the original curve functions β_1, β_2 . The \mathbb{L}^2 metric between β_1 and β_2 is not appropriate, as the action of re-parameterization is not by isometries: $(\int_0^1 |\beta_1(t) - \beta_2(t)|^2 dt)^{1/2} = (\int_0^1 |\beta_1(\gamma(t)) - \beta_2(\gamma(t))|^2 dt)^{1/2}$ does not hold for all valid re-parameterization functions γ (defined in the next section). However, the \mathbb{L}^2 metric with SRVFs is invariant under common re-parameterization. In addition, the metric involving SRVFs is related to the elastic metric [19, 5, 13], a Riemannian metric measuring bending and stretching energies required to deform one curve into another.

2.3. Locally Weighted Registration

One may note that Section 2.2 assumes that a given pair of curves q_1 and q_2 share common landmark locations along their respective arc-length parameterizations. This is generally not true: if $\theta_i^{(j)}$ is the i^{th} landmark on the j^{th} curve (for $j = 1, 2$), then $\theta_i^{(1)} \neq \theta_i^{(2)}$ in most cases. This means that the integration in Equations 2 and 3 is not well-defined due to the mismatch in integration bounds. In order to fix this, landmark constraints are initially matched by finding a suitable re-parameterization function γ_{lm} such that $\gamma_{\text{lm}}(\theta_i^{(1)}) = \theta_i^{(2)}$ for $i = 1, \dots, k$. This function should be an element of the group of re-parameterization functions $\Gamma = \{\gamma : [0, 1] \rightarrow [0, 1] | \gamma(0) = 0, \gamma(1) = 1, 0 < \dot{\gamma} < \infty\}$ (the set of all orientation-preserving diffeomorphisms of the unit interval). We choose to initialize with a piecewise linear function (with arc-length parameterization) between the matched landmarks. While this function is not a member of Γ (due to the required differentiability), it is generally a good approximation. If smoothness is desired, one can use the initialization step found in [17]. We define the landmark-matched second SRVF $q_{2,\text{lm}} := (q_2 \circ \gamma_{\text{lm}})\sqrt{\dot{\gamma}_{\text{lm}}}$.

While the landmarks are now matched, segments between landmarks are not. To align (or register) the shapes, we optimally match one shape to the other. This is done by minimizing the WLCP metric over the set of all valid rotations and re-parameterizations. We take into account rotation and re-parameterization invariance by defining the transformation groups of interest below. Let $SO(2)$ be the group of 2×2 rotation matrices. Since landmarks on both shapes are now matched, valid re-parameterization functions must fix these points. Thus, define the landmark-constrained re-parameterization subgroup $\Gamma_0 = \{\gamma \in \Gamma | \gamma(\theta_i) = \theta_i, i = 1, \dots, k\}$. The actions of O and Γ_0 commute (see [14] for proof): applying rotation $O \in SO(2)$ and re-parameterization $\gamma \in \Gamma_0$ to a SRVF q yields the transformed SRVF $O(q \circ \gamma)\sqrt{\dot{\gamma}}$. Define the *landmark-constrained size-and-shape space* $\mathcal{S}_s = \mathcal{C}_s / (SO(2) \times \Gamma_0)$.

This quotient space consists of equivalence classes $[q] = \{O(q \circ \gamma)\sqrt{\dot{\gamma}} | q \in \mathcal{C}_s, O \in SO(2), \gamma \in \Gamma_0\}$ which equate SRVFs that are only different by a rotation and/or re-parameterization. Thus each equivalence class uniquely identifies the object's size-and-shape.

To register q_1 and $q_{2,\text{lm}}$, we solve:

$$(O^*, \gamma^*) = \underset{O \in SO(2), \gamma \in \Gamma_0}{\operatorname{argmin}} d_{\mathcal{C}_s}^{(w)}(q_1, O(q_{2,\text{lm}} \circ \gamma)\sqrt{\dot{\gamma}}). \quad (4)$$

Solving Equation 4 is done by fixing q_1 , and searching for the rotation and re-parameterization which best matches $q_{2,\text{lm}}$, with respect to the WLCP metric. The solution can be approximated marginally, i.e., by finding O^* and γ^* separately and iterating until a stable solution is obtained.

2.3.1 Optimizing over $SO(2)$

Given a re-parameterization function $\tilde{\gamma}$, define $\tilde{q}_2 := (q_{2,\text{lm}} \circ \tilde{\gamma})\sqrt{\dot{\tilde{\gamma}}}$. The minimizing rotation is,

$$\begin{aligned} O^* &= \underset{O \in SO(2)}{\operatorname{argmin}} d_{\mathcal{C}_s}^{(w)}(q_1, O\tilde{q}_2) = \underset{O \in SO(2)}{\operatorname{argmax}} \langle \langle q_1, O\tilde{q}_2 \rangle \rangle^{(w)} \\ &= \underset{O \in SO(2)}{\operatorname{argmax}} \int_0^1 w(t) \langle q_1(t), O\tilde{q}_2(t) \rangle dt \\ &= \underset{O \in SO(2)}{\operatorname{argmax}} \operatorname{tr} \left(\left[\int_0^1 q_{1,w}(t) \tilde{q}_2(t)^\top dt \right] O^\top \right), \end{aligned} \quad (5)$$

where $q_{1,w}(t) := w(t)q_1(t)$. The last line of Equation 5 is solved by finding the singular value decomposition (SVD) of $A = \int_0^1 q_{1,w}(t) \tilde{q}_2(t)^\top dt = U\Sigma V^\top$, where the columns of U and V are singular vectors with corresponding singular values (in descending order) in diagonal matrix Σ . Then, the optimal rotation is given by $O^* = UV^\top$ (with the last column of V changing sign if $\det(A) = -1$). Note that this will differ from the optimal rotation under equal weights, as the SVD is impacted by the weight function $w(t)$ through the derived ‘‘weighted’’ SRVF $q_{1,w}$. By introducing variable weighting along the registered curves, the optimal rotation is dominated by the highest-weighted segments.

2.3.2 Optimizing over Γ_0

Given a rotation \bar{O} , define $\bar{q}_2 := \bar{O}q_{2,\text{lm}}$. The optimal re-parameterization is given by:

$$\begin{aligned} \gamma^* &= \underset{\gamma \in \Gamma_0}{\operatorname{argmin}} d_{\mathcal{C}_s}^{(w)}(q_1, (\bar{q}_2 \circ \gamma)\sqrt{\dot{\gamma}}) \\ &= \underset{\gamma \in \Gamma_0}{\operatorname{argmax}} \sum_{i=1}^{k+1} w_i \int_{S_i} \langle q_1(t), \bar{q}_2(\gamma(t))\sqrt{\dot{\gamma}(t)} \rangle dt. \end{aligned} \quad (6)$$




Figure 1. Curves β_1 and β_2 for Simulations 1–4 from Section 3 (size of curves not shown proportionally across simulations); landmark denoted by green circle.

The final expression in Equation 6 is linear along each curve piece, and thus can be optimized by maximizing each integral over S_i separately. This results in the weights dropping out of the optimization, meaning that optimizing over Γ_0 under the WLCP metric is equivalent to the same problem under the unweighted \mathbb{L}^2 metric. The solution can be obtained by employing a dynamic programming method for each piece separately, and “gluing” the solutions together (as discussed in [18]) or through gradient-descent [17], which ensures that the optimal re-parameterization is smooth everywhere (including at landmark constraints). Note that the optimal rotation and re-parameterization depend on each other; different weight functions will produce varying optimal rotations, which will impact the best correspondence of points between curves. Thus, in practice, it is suggested that one solves for the optimal rotation first before finding the optimal re-parameterization.

2.3.3 Weighted Metric on \mathcal{S}_s

Once the optimal rotation and re-parameterization pair (O^*, γ^*) is found (via Sections 2.3.1 and 2.3.2), let $q_2^* := O^*(q_{2,lm} \circ \gamma^*)\sqrt{\gamma^*}$ be the optimally-aligned second SRVF to the first SRVF. Since the action on $SO(2) \times \Gamma_0$ commutes and is by isometries on \mathcal{C}_s , \mathcal{S}_s inherits the WLCP metric on \mathcal{C}_s , and thus the *weighted landmark-constrained size-and-shape (WLCS) metric* between two size-and-shapes $[q_1], [q_2]$ is,

$$d_{\mathcal{S}_s}^{(w)}([q_1], [q_2]) = d_{\mathcal{C}_s}^{(w)}(q_1, q_2^*), \quad (7)$$

i.e., the minimizing distance over all possible rotations and re-parameterizations in \mathcal{C}_s . Note that introducing weights allows one to meticulously compare shapes according to local features. Upweighting certain segments allows for these

features to dominate registration. Then, the weighted distance calculated in Equation 7 will magnify any shape differences along the upweighted segments, allowing one to further scrutinize local shape differences.

2.4. Geodesic Paths

The WLCS metric can be interpreted as describing the amount of deformation required to go from one shape to the other (with respect to the selected weights). One can associate with this metric a locally shortest path in \mathcal{S}_s that shows how this deformation occurs. Mathematically, let $\alpha_{q_1, q_2} : [0, 1] \times [0, 1] \rightarrow \mathbb{R}^2$ be a path connecting q_1 to q_2 in \mathcal{C}_s (with the subscript sometimes dropped for convenience), i.e., $\alpha(0, t) = q_1(t)$ and $\alpha(1, t) = q_2(t)$, and let \mathcal{P} be the space of all paths from q_1 to q_2 in \mathcal{C}_s . Under the WLCP metric, the *geodesic path* $\hat{\alpha}_{q_1, q_2}^{(w)}$ is,

$$\begin{aligned} \hat{\alpha}_{q_1, q_2}^{(w)} &= \operatorname{argmin}_{\alpha \in \mathcal{P}} \int_0^1 \langle \dot{\alpha}(s, \cdot), \dot{\alpha}(s, \cdot) \rangle^{(w)} ds \\ &= \operatorname{argmin}_{\alpha \in \mathcal{P}} \int_0^1 \int_0^1 w(t) \langle \dot{\alpha}(s, t), \dot{\alpha}(s, t) \rangle dt ds \\ &= \operatorname{argmin}_{\alpha \in \mathcal{P}} \int_0^1 w(t) \left[\int_0^1 \langle \dot{\alpha}(s, t), \dot{\alpha}(s, t) \rangle ds \right] dt \\ &= \hat{\alpha}_{q_1, q_2}, \end{aligned} \quad (8)$$

where $\dot{\alpha}$ is the derivative of α with respect to the first argument s and $\hat{\alpha}_{q_1, q_2}(\cdot, t) = \operatorname{argmin}_{\alpha \in \mathcal{P}} \int_0^1 \langle \dot{\alpha}(s, t), \dot{\alpha}(s, t) \rangle ds$ is the geodesic path between $q_1(t)$ and $q_2(t)$ under the unweighted metric. We know that this path is linear, and so the geodesic path in \mathcal{C}_s is given by:

$$\hat{\alpha}_{q_1, q_2}^{(w)}(s, \cdot) = (1-s)q_1 + sq_2, \quad 0 \leq s \leq 1. \quad (9)$$

The corresponding geodesic path in \mathcal{S}_s is $\hat{\alpha}_{q_1, q_2}^{(w)}$, where q_2^* is the optimally aligned SRVF defined in Section 2.3.3. The WLCS metric given in Equation 7 is the *geodesic distance*, which is the length of the geodesic path. In practice, visualization of the geodesic path occurs by sampling G equally-spaced points along $\hat{\alpha}^{(w)}$, i.e., $\hat{\alpha}^{(w)}(0, \cdot), \hat{\alpha}^{(w)}(\frac{1}{G-1}, \cdot), \dots, \hat{\alpha}^{(w)}(\frac{G-2}{G-1}, \cdot), \hat{\alpha}^{(w)}(1, \cdot)$.

2.5. Closed Curves

Thus far, the weighted registration framework is defined for planar open curves. Closed curves impose an extra closure condition ($\int_{S^1} \dot{\beta}(t) dt = 0$), which results in a different shape space than for open curves. This occurs because, unlike open curves, closed curves do not have well-defined starting and ending points. However, landmark specification allows one to approximate closed curves as “unwrapped” open curves, where the first landmark θ_1 acts as a natural “starting point” for the curve β . Specifying k

Weights	Sim 1	Sim 2	Sim 3	Sim 4
[1.9, 0.1]	86.68	88.96	87.62	87.35
[1.5, 0.5]	71.33	83.43	78.18	71.92
[1, 1]	45.08	70.92	58.77	45.35
[0.5, 1.5]	18.86	43.94	29.23	18.77
[0.1, 1.9]	3.53	8.66	5.32	3.34

Table 1. Optimal clockwise rotation angle θ^* under different weight settings for Simulations 1–4 (all angles in degrees).

landmarks now splits a closed curve into k segments, which can be identified by partitioning the curve domain \mathbb{S}^1 into $S_1 = [\theta_1, \theta_2), S_2 = [\theta_2, \theta_3), \dots, S_{k-1} = [\theta_{k-1}, \theta_k), S_k = [\theta_k, \theta_1)$ (the last segment connects the final specified landmark back to the starting point). The curve domain \mathbb{S}^1 is commonly identified with the unit interval $[0, 1]$. The weight function defined in Section 2.2 is defined similarly to open curves, except with the constraint $\sum_i w_i = k$, and all registration procedures are identical. We approximate the geodesic path for closed curves using Equation 9 for open curves; however, in general, one can perform path straightening to find the actual solution computationally [13].

3. Simulated Curves

In this section, we illustrate the impact of weighting on registration of artificial curves. Figure 1 shows four pairs of two-peaked curves to be compared. All four simulations compare an original curve with both peaks oriented in the same way, to a second curve where one of the peaks is rotated by 90 degrees counterclockwise. In Simulation 1, all peaks are of equal size. At least one of the curves in each of Simulations 2–4 feature a peak which is three times higher than the other. In Simulation 2, the higher peak is simply rotated by 90 degrees to produce the second curve. Simulation 3 compares equally-sized peaks to a curve which has the higher peak rotated by 90 degrees. Lastly, Simulation 4 rotates the left peak by 90 degrees and switches the high and low peaks. For all examples, one landmark was specified to separate the two peaks, yielding two curve segments. Curves were not standardized to remove scale variability (in order to preserve the magnitude of peaks), so we are interested in comparing size-and-shapes. For all four simulations, registration was performed under the WLCP metric, under five different weight settings $[w_1, w_2]$ corresponding to the two curve segments: $[1.9, 0.1], [1.5, 0.5], [1, 1], [0.5, 1.5], [0.1, 1.9]$. The optimal rotation matrix $O^* = \begin{bmatrix} \cos(\theta^*) & \sin(\theta^*) \\ -\sin(\theta^*) & \cos(\theta^*) \end{bmatrix}$ (with optimal clockwise rotation angle θ^*) and re-parameterization function γ^* were found. Table 1 shows θ^* for all simulations under the five weight settings. Figure 2 shows the corresponding γ^* . Figure 3 shows geodesic paths under the various weights, with geodesic distances listed in Table 2.




Figure 2. Optimal re-parameterization γ^* for Simulations 1–4. Colors for different weight settings: black = $[1.9, 0.1]$, magenta = $[1.5, 0.5]$, blue = $[1, 1]$, red = $[0.5, 1.5]$, green = $[0.1, 1.9]$.

Weights	Sim 1	Sim 2	Sim 3	Sim 4
[1.9, 0.1]	0.573	0.565	1.574	1.694
[1.5, 0.5]	1.234	1.288	1.811	2.226
[1, 1]	1.468	1.756	1.954	2.460
[0.5, 1.5]	1.233	1.823	1.704	2.224
[0.1, 1.9]	0.581	1.035	0.842	1.690

Table 2. Geodesic distance $d_{S_s}^{(w)}([q_1], [q_2])$ under different weight settings for Simulations 1–4.

In Simulation 1, the optimal rotation is approximately 45° for equal weighting. This is clear, since neither peak dominates registration, meaning the best rotational alignment of the second curve is equally impacted by the perfect alignment of the second peak and the 90° rotation of the first peak. However, due to the symmetry of the matching, as the weight of one of the peaks is increased, the optimal rotation is pulled to match that particular peak (e.g., for $w_1 = 1.9$, the second shape is rotated by 86.68° to align with the first peak of the first shape). Also notice that the optimal re-parameterization function γ^* changes due to the weights. The blue line (with equal weighting) is roughly similar before and after the separation point of the two peaks. However, increasing a peak’s weight pushes γ^* to be almost linear for parameter values corresponding to that peak, and drastically changes it for the other peak to compensate. The geodesic paths found in Figure 3 suggest that deformation occurs differently depending on the weights as well. Constant weighting preserves peak structures equally well. Up-weighting one peak forces the deformation path to preserve

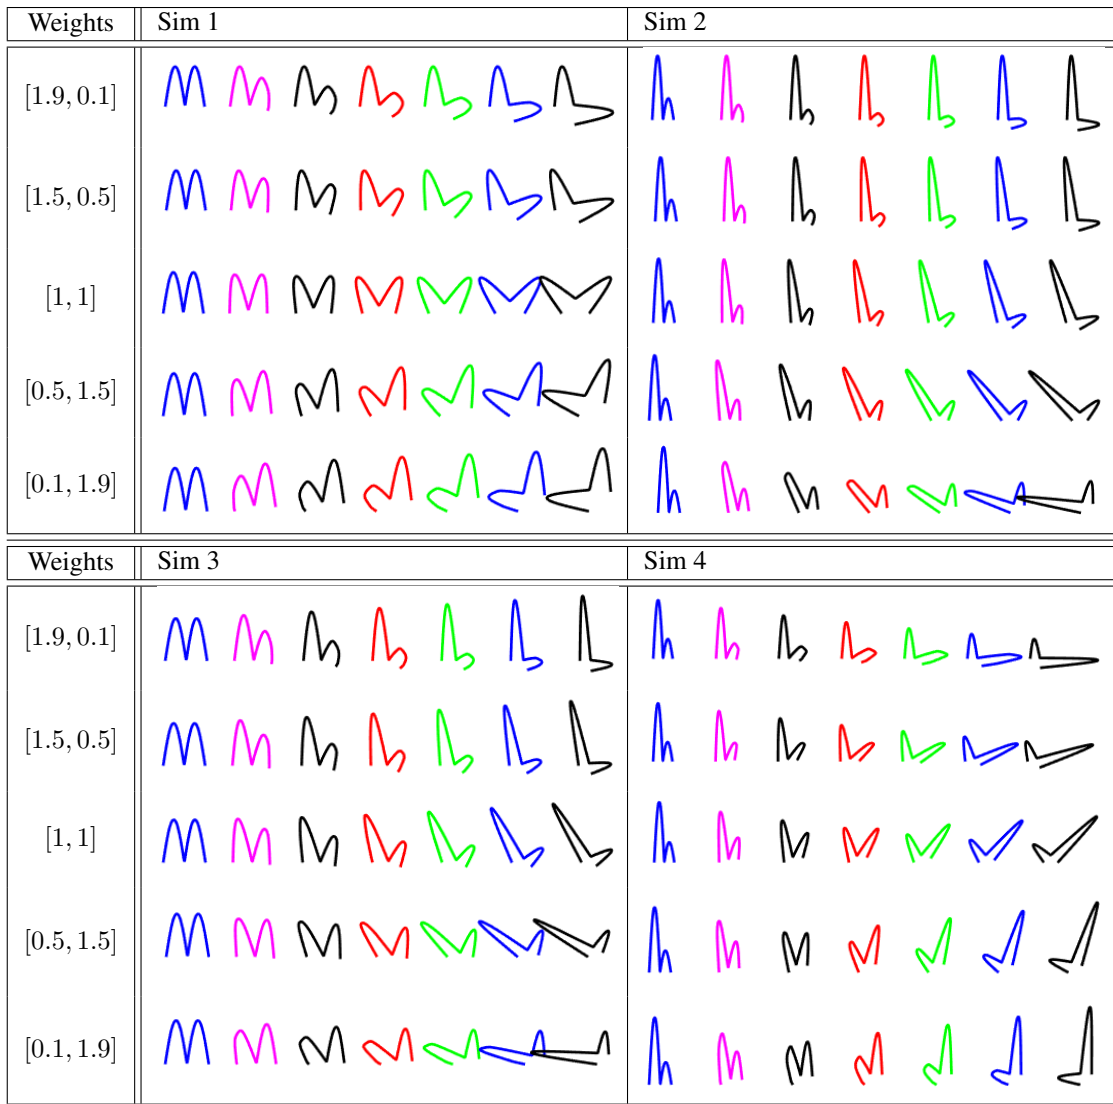


Figure 3. Geodesic paths for Simulations 1–4 under different weight settings; $G = 7$ sampled points along the path are displayed.

that particular peak as well as possible, while sacrificing the structure of the downweighted peak (e.g., notice how the second peak for $w_1 = 1.9$ shrinks and becomes asymmetric before growing and regaining symmetry towards the end).

Since the left peak is much larger in Simulation 2, this segment does not have to be weighted as strong as in Simulation 1 in order to optimally align this feature – $w_1 = 1.5$ recovers most of the rotation required to go from the first to second curve. However, in order to align with respect to the smaller feature, the segment must be upweighted strongly to overcome the first peak’s contribution to the optimal registration. Notice that again, γ^* varies by choice of weights. Finally, a result of note in Simulation 2 occurs when $w_1 = 0.5$. In all of the other simulations, the variable-weighted geodesic distance is smaller than the

equally-weighted geodesic distance, but in this case it is not. The optimal rotational alignment found does not match shapes very well at this setting, and so both peaks are quite different between the two curves. This raises an important point: the weighted metric is not guaranteed to be smaller than the unweighted metric, as the optimal rotation and re-parameterization functions will change! Thus, it does not make sense to compare distances across different weight settings, as changing weights influences registration.

In Simulations 3 and 4, peak sizes vary between the two curves. For Simulation 3, we again observe that the optimal rotation is mainly impacted by the dominant left peak in the second curve. As the left peak is upweighted, the optimal rotation will approach 90° ; however, this does not occur as “quickly” as in Simulation 2, since the first curve does

not have the same left-peak dominance. Notice the change point in the optimal re-parameterization which occurs at the point between the two peaks. This makes sense, since a larger portion of the second overall curve must be traversed in the same amount of time as the first half of the first curve. Also note that in the geodesic paths, the size of the first peak drastically changes in comparison to the second peak as the first peak is weighted higher; however, upweighting the second peak results in much slower evolution of the relative peak sizes. Simulation 4 alters the size of both the first and second peaks; optimal rotations reflect what is expected with the specified weights, and the geodesic paths display some symmetry (after rotation is factored out) with respect to weighting of the two peaks.

4. Examples

In this section, we perform pairwise weighted registration on objects from real data sets. As we are interested in shape (rather than size-and-shape), all objects are rescaled to lie on the unit Hilbert sphere \mathcal{C} . However, the extrinsic WLCS metric (Equation 7) is still used for comparison.

4.1. MPEG-7 Shapes from Computer Vision

We first demonstrate the impact of weighted registration on shapes from the MPEG-7 data set¹, which features complex shapes from computer vision. Figure 4 shows two sets of shapes for comparison: bones, with one “healthy” and one fractured bone, and half-circles, where there is approximate symmetry in one versus asymmetry in the other. In the bone example, $k = 6$ landmarks were selected, as shown at the top of the figure. Registration and geodesic calculation was done using three weight settings: (1) $[1, 1, 1, 1, 1, 1]$ (equal weights), (2) $[1.9, 1.9, 0.1, 0.1, 0.1, 1.9]$ (right side), and (3) $[0.1, 0.1, 1.9, 1.9, 1.9, 0.1]$ (fractured left side). Notice that while the equal weight geodesic path looks like a natural deformation, perhaps it makes more sense to register with respect to the healthy part of the bone or the fracture. In real settings, this may allow a doctor to diagnose the severity of the fracture with more confidence. The rotational alignment differs significantly depending on which bone segment is upweighted (and this is also reflected in the optimal re-parameterizations). Notice that (2) seems to fix the healthy part of the bone and only slowly bend the left part to create the bend fracture. Optimal rotation angles are $\theta^* = 0.66^\circ, 9.53^\circ,$ and 15.41° , respectively. For the half-circle, $k = 3$ landmarks were selected and weight settings were as follows: (1) $[1, 1, 1]$, (2) $[2.8, 0.1, 0.1]$, and (3) $[0.1, 2.8, 0.1]$. The equally-weighted geodesic shows some rotational alignment in order to match the asymmetry of the left side more closely, while sacrificing the flat segment at the base. However, if the flat segment is upweighted

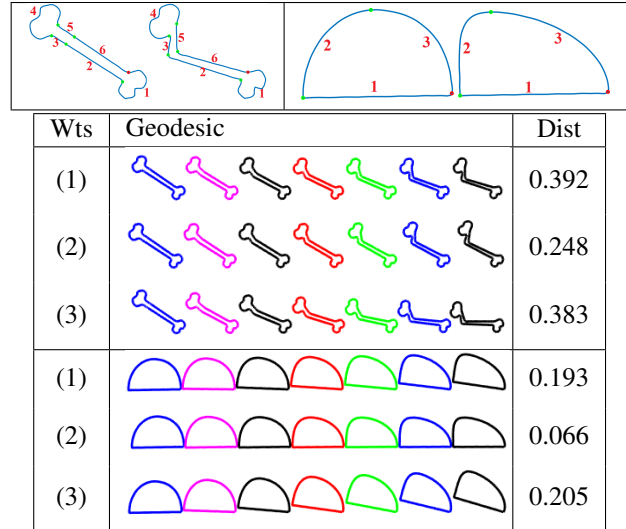


Figure 4. Top: Bone and half-circle shapes being compared – segments between landmarks are labeled. Bottom: Geodesic path and distance for different weight settings.

(as in (2)), this feature is preserved and the asymmetry is formed with respect to this feature. (3) upweights the left side, which forces a stronger rotational alignment in order to match the curvature of this segment as much as possible. Optimal rotation angles for (1), (2), and (3) are $\theta^* = 9.18^\circ, 0.76^\circ,$ and 16.97° , respectively. The left panel of Figure 6 shows optimal re-parameterizations for this example.

4.2. Mice Vertebrae

This section looks at shape of the second thoracic vertebra of mice from the ‘shapes’ package in R, developed and referenced by [4]. Further discussion of vertebra anatomy and grouping of mice can be found in [18]. A comparison of two mice vertebrae is shown in Figure 5, with $k = 2$ landmarks specified to separate the neural spine (the “tail” on the right side) from the rest of the structure. Weighted registration was performed using weights (1) $[1, 1]$, (2) $[1.9, 0.1]$ (ignoring neural spine), and (3) $[0.1, 1.9]$ (favoring neural spine alignment). Notice that in cases (1) and (2), since the neural spine is a small-scale feature as compared to the full shape, the optimal rotation is found to match the rest of the vertebrae ($\theta^* = 19.25^\circ$ and 21.87° , respectively). This seems to ignore the tilt difference in the neural spine. However, by upweighting the neural spine, we are able to align based on this local shape feature ($\theta^* = 9.35^\circ$), allowing the researcher to compare the vertebrae most predominantly using the neural spine. By aligning according to this feature, one could potentially form clusters of mice vertebrae based on the neural spine’s tilt, allowing for improved symmetry analysis of the structure.

¹<http://www.dabi.temple.edu/shape/MPEG7/dataset.html>

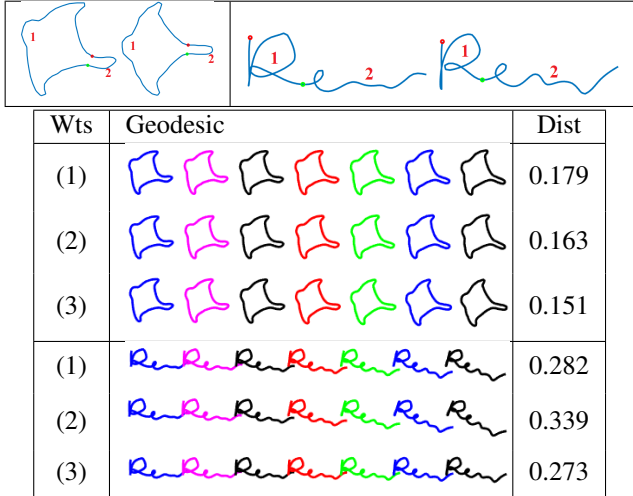


Figure 5. Top: Mice vertebrae and signatures being compared – segments between landmarks are labeled. Bottom: Geodesic path and distance for different weight settings.

4.3. Handwriting Samples

Figure 5 also shows a comparison of a signature for the name ‘Ren.’ This type of data is relevant in forensic analysis, where the goal is to discriminate between legitimate signatures and forgeries. The data used here come from the SVC 2004 data set², which contains 40 samples of 40 signatures each (including legitimate ones and forgeries) – see [8] for further details. Signatures can benefit from weighted registration, as in the unweighted case, registration may be dominated by a single large letter (e.g., the first one in a name); this could lead to misalignment of the full signature. By downweighting the first letter, we can align based on the majority of the letters in the signature and scrutinize those letters more carefully. For this example, the first signature is legitimate, while the second is a forgery. One landmark was placed at the end of the letter ‘R’, and we compute optimal registrations and geodesics under weight settings (1) [1, 1], (2) [1.9, 0.1] (emphasizing ‘R’ matching), and (3) [0.1, 1.9] (emphasizing ‘en’ matching). Notice that in the unweighted case (1), the rotational alignment appears off ($\theta^* = 14.11^\circ$) because the first letter dominates. As expected, upweighting the first letter will induce a stronger rotation ($\theta^* = 21.35^\circ$) to improve that letter’s alignment even more. However, downweighting the first letter (as in (3)) corrects for the overcompensation due to the first letter, and better aligns the rest of the signature ($\theta^* = 5.60^\circ$). The right panel of Figure 6 shows optimal re-parameterization functions for the weight settings; notice the similarities since the optimal rotational alignments are fairly similar across weights. Even when accounting for variability in orientation due to

²<http://www.cse.ust.hk/svc2004/download.html>




Figure 6. Optimal re-parameterization γ^* for half-circle (L) and signature (R) examples. Colors for different weight settings, respectively: magenta = [2.8, 0.1, 0.1] (L)/[1.9, 0.1] (R), blue = [1, 1, 1] (L)/[1, 1] (R), black = [0.1, 2.8, 0.1] (L)/[0.1, 1.9] (R).

the first letter, we still obtain a relatively large distance, indicating that the two signatures are different. Further scrutiny of individual letters in a signature can be done by manually placing landmark(s) to segment the particular letter of interest.

5. Summary and Future Work

In this work, we have demonstrated the ability to register shapes via local features by introducing a weight function. The optimal registration depends on selected landmarks and weights. By finding the optimal alignment, one can then compute pairwise distances between shapes (or size-and-shapes) using a weighted \mathbb{L}^2 metric between SRVFs. This metric also minimizes the registration energy, and produces a weighted geodesic path which illustrates shape deformations. For shapes with significant features that are rotationally misaligned, introducing weights can alter the optimal rotation by a large amount, which then impacts the optimal re-parameterization function. We demonstrated this on simulated curves where rotational misalignments were artificially introduced. Noticeable differences are also found in real shape data.

As this is an introduction to weighted registration, there are ample opportunities for future work. A natural concern is of choice of weights. Alternatively, one could learn the weights, given selected landmarks, by optimizing over a task of interest (e.g., classification performance). By doing this, the researcher can infer the most important local features as those with the highest weights. Establishing a weighted metric opens up the ability to look at mean calculation, variability assessment, and further inference under this new metric, and examine robustness to the choice of weights. Notice that in some examples, certain weight settings do not significantly impact registration. If any selection of weight yields a similar optimal rotation and alignment, then this could be thought of as a way of assessing heterogeneity in a population of shapes. The seamless incorporation of automatic landmark detection (in cases where landmarks are not pre-specified) for classes of similar shapes is also an important future direction for this work.

References

- [1] M. Bauer, M. Eslitzbichler, and M. Grasmair. Landmark-guided elastic shape analysis of human character motions. *Inverse Problems and Imaging*, 11(4):601–621, 2017. [1](#)
- [2] F. L. Bookstein. Size and shape spaces for landmark data in two dimensions. *Statistical Science*, 1(2):181–222, 1986. [1](#)
- [3] K. Domijan and S. P. Wilson. A Bayesian method for automatic landmark detection in segmented images. In *International Conference on Machine Learning*, 2005. [2](#)
- [4] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis: with Applications in R, Second Edition*. Wiley, New York, 2016. [1](#), [7](#)
- [5] S. H. Joshi, E. Klassen, A. Srivastava, and I. H. Jermyn. A novel representation for Riemannian analysis of elastic curves in r^n . In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7, 2007. [3](#)
- [6] D. G. Kendall. Shape manifolds, Procrustean metrics, and complex projective shapes. *Bulletin of London Mathematical Society*, 16:81–121, 1984. [1](#)
- [7] E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(3):372–383, 2004. [1](#)
- [8] S. Kurtek and A. Srivastava. Handwritten text segmentation using elastic shape analysis. In *International Conference on Pattern Recognition*, 2014. [8](#)
- [9] W. Liu, A. Srivastava, and J. Zhang. Protein structure alignment using elastic shape analysis. In *ACM International Conference on Bioinformatics and Computational Biology*, 2010. [1](#)
- [10] P. W. Michor, D. Mumford, J. Shah, and L. Younes. A metric on shape space with explicit geodesics. *Matematica E Applicazioni*, 19:25–57, 2007. [1](#)
- [11] C. Prematilake and L. Ellingson. Evaluation and prediction of polygon approximations of planar contours for shape analysis. *Journal of Applied Statistics*, pages 1–20, 2017. [2](#)
- [12] C. G. Small. *The Statistical Theory of Shape*. Springer, 1996. [1](#)
- [13] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn. Shape analysis of elastic curves in Euclidean spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33:1415–1428, 2011. [1](#), [2](#), [3](#), [5](#)
- [14] A. Srivastava and E. P. Klassen. *Functional and Shape Data Analysis*. Springer-Verlag, 2016. [3](#)
- [15] J. Strait, O. Chkrebti, and S. Kurtek. Automatic landmark detection and uncertainty quantification of landmarks on elastic curves. *arXiv:1710.05008*, 2017. [2](#)
- [16] J. Strait and S. Kurtek. Bayesian model-based automatic landmark detection for planar curves. In *IEEE CVPR Workshop on Differential Geometry in Computer Vision and Machine Learning*, 2016. [2](#)
- [17] J. Strait and S. Kurtek. A novel algorithm for optimal matching of elastic shapes with landmark constraints. In *International Conference on Image Processing Theory, Tools, and Applications*, 2017. [3](#), [4](#)
- [18] J. Strait, S. Kurtek, E. Bartha, and S. MacEachern. Landmark-constrained elastic shape analysis of planar curves. *Journal of the American Statistical Association*, 112(518):521–533, 2017. [1](#), [2](#), [4](#), [7](#)
- [19] L. Younes. Computable elastic distance between shapes. *SIAM Journal of Applied Mathematics*, 58(2):565–586, 1998. [1](#), [3](#)
- [20] C. T. Zahn and R. Z. Roskies. Fourier descriptors for plane closed curves. *IEEE Transactions on Computers*, 21(3):269–281, 1972. [1](#)