A Novel Algorithm for Optimal Matching of Elastic Shapes with Landmark Constraints

Justin Strait and Sebastian Kurtek
The Ohio State University
e-mail: strait.50@osu.edu, kurtek.1@stat.osu.edu

***Abstract*— An important problem in statistical shape analysis is the matching of geometric features across shapes, known as registration. In short, given two objects, one wants to know the correspondence of points on one shape to points on another. Such a matching problem, with various levels of complexity, is present regardless of the shape's mathematical representation. A recent framework for shape analysis of n-dimensional curves combines an infinite-dimensional functional curve representation with landmark information encoding important curve features. In this setting, shape matching is performed by minimizing an objective function with constraints, which respect landmark correspondences. Currently, the minimizer in this approach is found using piecewise dynamic programming; this does not respect the smoothness requirement of the matching function. Thus, the solution is not really a member of the group of registration functions. In this work, we present a landmark-constrained gradient descent algorithm, which searches for a smooth matching function and respects landmark locations. We compare the proposed method to the previously used approach using examples from the MPEG-7 dataset.**

***Keywords*— shape analysis, landmarks, SRVF, registration**

I. Introduction

The field of statistical shape analysis seeks to define tools for representation, registration (matching), comparison and various statistical analysis tasks for shape data. This type of data is now common in many fields including biometrics, biology, medicine, graphics, bioinformatics, etc. For instance, researchers are often interested in correlating the shape of outlines of anatomical structures in medical images to various diseases. One major challenge in statistical shape analysis is the choice of shape representation. Early literature on this subject represented objects using a finite set of labeled points called landmarks [7]. The shape of an object was then defined as a property that is unaffected by the translation, rotation and scaling of the landmark configuration. A lot of sophisticated statistical machinery for shape analysis was developed under this representation [4]. The main drawback of landmark-based shape analysis is that only few points are used to represent the entire object, i.e., there is significant loss of information. Mathematical and technological advances have led to the development of functional shape representations, where the entire object boundary is treated as the data [16], [6], [13]. These approaches prove more challenging, both conceptually and computationally, due to an additional invariance requirement of shapes: a reparameterization of the function representing the object of interest does not alter its shape. All shape representations must account for the desired invariances before statistical procedures can be employed.

The most basic task in statistical shape analysis is comparison, which requires the definition of a metric on an appropriate shape space. This distance must be preserved under all invariances listed previously. For functional shape analysis approaches, this typically relies on finding a set of optimal transformations, which provide a matching of one shape to the other. In practice, this is most commonly done by searching for a translation, rescaling, rotation and reparameterization which are distance-minimizing. An optimal reparameterization seeks to find the best "matching" of every point on one curve to points on the other. The first three transformations can be found rather simply using existing Procrustes methods. The last transformation is the hardest to search for as it involves optimizing over the infinite-dimensional group of diffeomorphisms (we make the formulation of this problem more precise in later sections).

This work addresses the search for an optimal reparameterization within the landmark-constrained elastic shape analysis framework [15]. Our method is based on an efficient landmark-constrained gradient descent algorithm, and the proposed formulation guarantees a solution in the correct space, i.e., the space of diffeomorphisms.

A. Previous Work

Gradient descent algorithms for optimal reparameterization searches have been previously developed in the context of elastic shape analysis when no landmark constraints are present. In particular, Srivastava *et al.* [13], [14] describe one such algorithm and derive all necessary components needed to implement it. The algorithm proposed here is similar in spirit to this previous work, albeit with the additional necessity of preserving landmark correspondences. This development poses new challenges in the specification of the energy gradient as described later. A very common tool for unconstrained registration of time indexed curves is the dynamic programming (DP) algorithm, also commonly referred to as dynamic time warping (DTW) [3]. Unfortunately, as shown in [9], [13], the naive DTW approach does not lead to a proper metric on the shape space of curves. Thus, for landmark-free statistical shape analysis, a modified registration problem, which does not suffer from the same drawbacks, is usually defined. Then, a modified version of the DP approach is used to solve this problem; details can be found in [12], [9], [13]. Avant and

978-1-5386-1842-4/17/$31.00

Gee [1] consider a similar problem of matching curves with constraints under a morphometric framework.

Another popular shape matching approach is called Large Deformation Diffeomorphic Metric Mapping (LDDMM) [2], [5]. While these methods have considered both landmark-based and functional shape representations, they solve an entirely different problem. The formulation in those frameworks seeks a diffeomorphism of the ambient space (e.g., $\mathbb{R}^2$ for planar curves) such that the two objects are optimally deformed into each other. In the current work, we seek a landmark-constrained diffeomorphism of the curve domain (either $[0,1]$ or $\mathbb{S}^1$) such that points across the two curves are optimally matched. In fact, there is very little existing literature on the topic of landmark-constrained registration of curves under a functional representation. One recent example considers a piecewise dynamic programming (PDP) approach [15]; in short, the given curves are first segmented into multiple open curves based on the landmark constraints, and then a separate DP algorithm is used to match each pair of these segments. Unfortunately, as shown later, such an approach does not generally result in a smooth reparameterization function. Our work is most related to that of Kurtek *et al.* [10], which considered landmark-constrained elastic shape analysis of closed surfaces.

The rest of this paper is organized as follows. Section II describes the optimization problem of interest and the proposed solution. In Section III, we demonstrate our approach on complex examples, and compare it to the previously used PDP method. Finally, we close with a brief summary and some directions for future work in Section IV.

II. Landmark-Constrained Shape Registration

In this section, we define the constrained optimization problem of interest and derive the proposed gradient descent solution. This includes a description of an appropriate initialization scheme as well as a subsequent landmark-constrained search over a subgroup of diffeomorphisms. Throughout this paper we use $|\cdot|$ and $\langle\cdot,\cdot\rangle$ to denote the Euclidean norm and inner product in $\mathbb{R}^n$, respectively; we also use $\|\cdot\|$ and $\langle\langle\cdot,\cdot\rangle\rangle$ to denote the $\mathbb{L}^2$ norm and inner product, respectively.

A. Preliminaries

Let $\beta_i : D \to \mathbb{R}^2,\ i = 1,2$ be two absolutely continuous, planar curves with domain D (for open curves $D = [0,1]$ and for closed curves $D = \mathbb{S}^1$). In this paper, we only consider planar curves, but the proposed method applies to n-dimensional curves. In order to compare shapes of β_1 and β_2, we require a metric which respects the invariances imposed by our definition of shape; specifically, it should be invariant to the parameterizations of the curves (in addition to the standard similarity group). We restrict the following discussion to open curves. A reparameterization function is a diffeomorphism of $[0,1]$, i.e., it is a member of the following Lie group: $\Gamma = \{\gamma : [0,1] \to [0,1] | \gamma(0) = 0,\ \gamma(1) = 1,\ \gamma \text{ is an orientation-preserving diffeomorphism}\}$. Note that group membership requires an increasing, smooth function. Elements $\gamma \in \Gamma$ act on a curve β by composition: $(\beta, \gamma) = \beta \circ \gamma$. The tangent space at the identity element, $\gamma_{id}(t) = t$, represents all possible "perturbations" of γ_{id}. Formally, this tangent space contains smooth functions, which vanish at the boundaries: $T_{\gamma_{id}}(\Gamma) = \{v : [0,1] \to \mathbb{R} | v(0) = v(1) = 0,\ v \text{ is smooth}\}$.

It can be shown that the $\mathbb{L}^2$ metric on the space of absolutely continuous curves is not preserved under reparameterizations; this implies that any shape analysis framework, which uses this metric to compare shapes, is not parameterization invariant. However, Srivastava *et al.* [13] show that the $\mathbb{L}^2$ metric on the space of *square-root velocity function* (SRVF) transformations of absolutely continuous curves is invariant to reparameterizations; the SRVF is defined as $q = \frac{\dot{\beta}}{\sqrt{|\dot{\beta}|}}$, where $\dot{\beta}$ is the time-derivative of β (the inverse mapping is also simple to compute). Under this representation, an element $\gamma \in \Gamma$ acts on an SRVF q as $(q, \gamma) = (q \circ \gamma)\sqrt{\dot{\gamma}}$. The primary benefit of the SRVF representation is the ability to use the $\mathbb{L}^2$ metric to find optimal reparameterizations of curves; it also corresponds to an elastic metric on the original space of curves, which measures the amount of bending and stretching required to deform one shape into another.

B. Landmark Constraints

Given a metric, one can find an optimal reparameterization between two planar curves by minimizing the distance between them. However, optimizing over elements in Γ does not preserve landmark locations; recall that landmarks are selected to be in correspondence across shapes. Thus, to constrain the problem, Strait *et al.* [15] define a subgroup of Γ which respects landmark locations. Suppose that each curve is annotated using k labeled landmark points $\beta_i(\mathrm{s}) = (\beta_i(s_1), \ldots, \beta_i(s_k)),\ i = 1,2$, where $s_j \in [0,1]$ for all j. The landmark-constrained reparameterization group is defined as $\Gamma_0 = \{\gamma : [0,1] \to [0,1] | \gamma(0) = 0,\ \gamma(1) = 1,\ \gamma(s_j) = s_j,\ j = 1, \ldots, k,\ \gamma \text{ is an orientation-preserving diffeomorphism}\}$. Elements of Γ_0 preserve landmark correspondences across shapes. Because Γ_0 is a subgroup of Γ, the group action remains the same.

The optimal landmark-constrained reparameterization is then defined as the solution to the following minimization problem:

$$\gamma^* = \underset{\gamma \in \Gamma_0}{\operatorname{argmin}}\ H(\gamma; q_1, q_2) = \underset{\gamma \in \Gamma_0}{\operatorname{argmin}} \|q_1 - (q_2, \gamma)\|^2 . \qquad (1)$$

DP efficiently provides the optimal solution to this problem when the search is over Γ rather than Γ_0. Strait *et al.* [15] find the optimal solution over Γ_0 by employing a PDP algorithm, i.e., instead of directly searching for a landmark-constrained diffeomorphism of $[0,1]$, they find the optimal matching for each pair of curve segments after landmark-based segmentation. In that case, the start and end points of each piecewise matching function must correspond to a landmark location, thus reducing the problem to multiple independent

DP searches for a diffeomorphism of $[a,b] \subset [0,1]$. While that approach provides an efficient, approximate solution to the problem at hand, the resulting γ is not smooth at the landmark locations, and thus, is not a member of Γ_0. To alleviate this problem, we propose an alternative gradient descent method which searches directly over Γ_0. We refer to this approach as landmark-constrained gradient descent (LCGD).

1) Initialization: The first step of the proposed algorithm is to match the k given landmarks across the two curves. That is, we seek a $\gamma^0 \in \Gamma$, which solves the following minimization problem:

$$\gamma^0 = \underset{\gamma\in\Gamma}{\operatorname{argmin}}\, E(\gamma; q_1, q_2) = \underset{\gamma\in\Gamma}{\operatorname{argmin}}\, |\gamma(\mathrm{s}^1) - \mathrm{s}^2|^2 + \lambda \left\|\gamma - \gamma_{int}\right\|^2, \tag{2}$$

where s^i denotes the set of landmark locations on curve i. To prevent the minimization from being ill-posed, the second term is included as a regularizer to push the solution toward γ_{int} (as the parameter λ is increased), the reparameterization function which satisfies $\gamma_{int}(\mathrm{s}^1) = \mathrm{s}^2$ and is parameterized by arc-length between landmark constraints. Since landmarks are defined as important points on objects that are in correspondence across a population of shapes, this initialization step searches for an initial reparameterization which only matches the landmark sets. Such a function always exists in Γ, and we find it using gradient descent.

First, we calculate the gradient of E, which can be approximated using an orthonormal basis of $T_{\gamma_{id}}(\Gamma)$. It can be shown that the directional derivative of E in the direction of $v \in T_{\gamma_{id}}(\Gamma)$ is:

$$\frac{d}{d\epsilon} E(\gamma + \epsilon v)\big|_{\epsilon=0} = 2\langle v(\mathrm{s}^1), \gamma(\mathrm{s}^1) - \mathrm{s}^2\rangle + 2\lambda\langle\langle \gamma - \gamma_{int}, v\rangle\rangle. \tag{3}$$

A basis of $T_{\gamma_{id}}(\Gamma)$ is given by [14]:

$$\left\{ \frac{1}{\sqrt{2}\pi n}\sin(2\pi n t),\ \frac{1}{\sqrt{2}\pi n}(\cos(2\pi n t) - 1) | n = 1, 2, \ldots \right\}. \tag{4}$$

This basis is orthonormal under the first-order Palais metric [11]. Denoting these basis elements by the set $\{b_i, i = 1, 2, \ldots\}$, the full gradient of E can be approximated by truncating this basis and calculating the quantity:

$$\nabla E_{\gamma_{id}}(b) \approx \sum_{i=1}^{n} \Big(\langle b_i(\mathrm{s}^1), \gamma(\mathrm{s}^1) - \mathrm{s}^2\rangle + \lambda\langle\langle \gamma - \gamma_{int}, b_i\rangle\rangle \Big) b_i, \tag{5}$$

for large n. This leads to Algorithm 1, which solves the initialization problem.

Algorithm 1: Initialization

1) Initialize $k = 0$, $\gamma_k = \gamma_{id}$, $\epsilon_1 > 0$ and $\delta_1 > 0$.
2) Compute $\nabla E_{\gamma_{id}}(b)$ using Equation 5 and γ_k.
3) Update by setting $\gamma_{[k+1]} = \gamma_{id} - \epsilon_1 \nabla E_{\gamma_{id}}(b)$ and compute $\gamma_{k+1} = \gamma_k \circ \gamma_{[k+1]}$.
4) If $E(\gamma_{k+1}) < \delta_1$, return $\gamma^0 = \gamma_{k+1}$. Otherwise, set $k = k + 1$ and return to step 2.

Through several simulations, we confirmed that the overall LCGD framework is robust to the choice of the step size ϵ_1 and the regularization parameter λ; we note that for a sufficiently small ϵ_1, this algorithm guarantees to search in Γ only. We give our choices of these parameters in Section III, which appear to work well for all of the experiments.

2) Landmark-Constrained Optimization: The second step of the proposed algorithm is to find γ^*, which solves the minimization problem in Equation 1, given the initial value $\gamma^0(\mathrm{s}^1) \approx \mathrm{s}^2$. This optimization is solved using gradient descent under a landmark-constrained basis. We provide the details of the algorithm next.

Let $\eta(\gamma) = (q \circ \gamma)\sqrt{\dot{\gamma}}$. The gradient of H can be approximated using directional derivatives of H given by (in any direction $v \in T_{\gamma_{id}}(\Gamma_0)$):

$$\frac{d}{d\epsilon} H(\gamma + \epsilon v)\big|_{\epsilon=0} = -2\langle\langle q_1 - (q_2, \gamma), \eta_*(\gamma)\rangle\rangle, \tag{6}$$

where $\eta_* = \frac{d}{d\epsilon}\eta(\gamma_{id} + \epsilon v)\big|_{\epsilon=0} = \dot{q}v + \frac{1}{2}q\dot{v}$ is the differential of η. In order to find the gradient of H, we require an orthonormal basis for $T_{\gamma_{id}}(\Gamma_0)$ defined as $T_{\gamma_{id}}(\Gamma_0) = \{v : [0,1] \to \mathbb{R} | v(0) = v(s_j) = v(1) = 0,\ j = 1, \ldots, k,\ v \text{ is smooth}\}$.

Unfortunately, the basis used for $T_{\gamma_{id}}(\Gamma)$ from the initialization step is not appropriate, as those functions do not vanish at landmark locations. To construct an appropriate basis, we begin with the basis in Equation 4 and truncate it to result in $\mathcal{B} = \{b_i,\ i = 1, \ldots, n\}$. To make sure that the basis functions vanish at the landmarks, we define a new basis set as follows:

1) Let $s_1, \ldots, s_k$ be the landmark locations and $j = 1$.
2) Choose the first $b_l \in \mathcal{B}$ where $b_l(s_j) \neq 0$.
3) For all $i \neq l$, set $\tilde{b}_i = b_i - \frac{b_i(s_j)}{b_l(s_j)} b_l$.
4) Remove b_l from $\mathcal{B}$ and set $j = j+1$. Repeat until $j = k$.

The set $\widetilde{\mathcal{B}} = \{\tilde{b}_i,\ i = 1, \ldots, n-k\}$ forms an approximate basis for $T_{\gamma_{id}}(\Gamma_0)$. This basis is orthonormalized using the Gram-Schmidt procedure under the first-order Palais metric. Then, we can approximate the full negative gradient of H by:

$$-\nabla H_{\gamma_{id}}(\tilde{b}) \approx \sum_{i=1}^{n-k} \langle\langle q_1 - q_2, \dot{q}_2 \tilde{b}_i + \frac{1}{2} q_2 \dot{\tilde{b}}_i\rangle\rangle \tilde{b}_i. \tag{7}$$

Note that while Γ_0 is not closed (i.e., the limit is not in the group), this is not an issue because the closure will never be obtained due to basis truncation. Algorithm 2 summarizes the steps required to find the optimal landmark-constrained reparameterization, which provides a correspondence between two curves β_1 and β_2.

Algorithm 2: Landmark-constrained reparameterization

1) Set $k = 0$, $\gamma_k = \gamma^0$ from Algorithm 1, $\epsilon_2 > 0$ and $\delta_2 > 0$.
2) Compute $\tilde{q}_2 = (q_2, \gamma_k)$.
3) Compute $\nabla H_{\gamma_{id}}(\tilde{b})$ using Equation 7 and γ_k.
4) Update by setting $\gamma_{[k+1]} = \gamma_{id} - \epsilon_2 \nabla H_{\gamma_{id}}(\tilde{b})$ and compute $\gamma_{k+1} = \gamma_k \circ \gamma_{[k+1]}$.
5) If $\left\|\nabla H_{\gamma_{id}}(\tilde{b})\right\| < \delta_2$, return $\gamma^* = \gamma_{k+1}$. Otherwise, set $k = k + 1$ and return to step 2.

In Algorithm 2, γ_k represents the overall reparameterization, which is updated by composition using small perturbations of

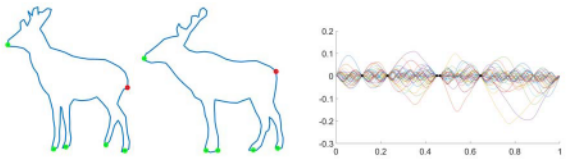

Fig. 1. Left: Outlines of two deer. The starting point is marked in red, and the remaining landmarks in green. Right: Approximate landmark-constrained basis (truncated to 35 elements). Black points represent landmark locations.

γ_{id} via the incremental functions $\gamma_{[k+1]}$; these are found by moving in the direction of the negative gradient in $T_{\gamma_{id}}(\Gamma_0)$. Again, the step size ϵ_2 is chosen appropriately for fast convergence of the algorithm.

III. Experimental Results

In this section, we demonstrate the gradient-based optimization over Γ_0 on complex shapes from the MPEG-7 dataset[1]. For comparison, we also compute optimal reparameterizations using the PDP approach used by Strait *et al.* [15].

Motivating Example. An important use of landmark-constrained shape analysis is in the comparison of two objects which contain occluded features. This is well illustrated using the outlines of two deer from the MPEG-7 dataset; they are shown in Figure 1. The main difference between the two deer is the occlusion of one of the front legs for the deer on the right. As shown in [15], unconstrained elastic shape analysis does not recognize the presence of the hidden front leg, leading to an incorrect matching of features and an unnatural deformation between the two shapes. Incorporating landmark constraints allows the analysis to be guided by this additional semantic information.

To compare the proposed LCGD registration method to PDP, we mark six total landmarks on each deer (marked in green and red); for closed curves, the first landmark (red) is used to identify $t = 0$. The deer with an occluded leg is carefully marked with two landmarks near the occlusion. We implement this comparison using $N = 600$ points on each curve (specifying 100 points for each segment between landmarks). Both deer are converted to their corresponding SRVF representations and re-scaled to unit length. The optimal rotation of the second deer to the first deer is also found prior to finding the optimal reparameterization.

To find the optimal landmark-constrained reparameterization to match the two deer using the LCGD method, we first use Algorithm 1 to match the given landmarks. We specify $n = 40$ basis elements to calculate the gradient given in Equation 5, and set $\epsilon_1 = 0.1$, $\delta_1 = 1 \times 10^{-5}$, and $\lambda = 1$. Once this initialization step is completed, we then optimize over Γ_0 using Algorithm 2. Calculating the gradient given in Equation 7 requires the landmark-constrained basis, which vanishes at the landmark locations. For closed curves, the starting point is always identified with $t = 0$ where all of the original basis elements already vanish. We additionally

[1] http://www.dabi.temple.edu/~shape/MPEG7/dataset.html

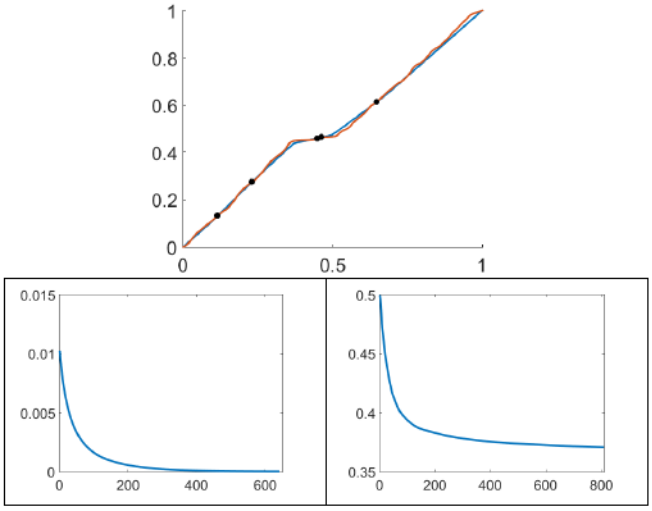

Fig. 2. Top: Initialization solution in blue and final LCGD solution in red. Black points mark landmark locations. Bottom: Energy for the initialization step (left) and landmark-constrained optimization step (right).

apply the previously described procedure to ensure that each basis function vanishes at the other five landmarks. The resulting landmark-constrained orthonormal basis, $\widetilde{\mathcal{B}}$, contains 35 elements and is shown in Figure 1. All of the displayed functions vanish at the landmark locations (shown using black points), which verifies the previously stated requirement. It is clear that these 35 basis elements should be sufficient for the subsequent optimization, i.e., both large and small-scale changes to the initialized reparameterization are available. We specify $\epsilon_2 = 0.005$ and $\delta_2 = 0.01$ in Algorithm 2. Unless stated otherwise, we use the same parameter settings for Algorithms 1 and 2 throughout this section.

The resulting initialization (blue) and optimal landmark-constrained reparameterization (red) solutions to Equations 2 and 1, respectively, are shown in Figure 2; also shown are energies E and H, respectively, as functions of the iteration number. The initialization required 638 iterations to converge; the subsequent optimization over Γ_0 required 805 additional iterations (note that these two numbers depend on the values of ϵ_1 and ϵ_2). Since the optimal reparameterization is obtained using the negative gradient of H with respect to the landmark-constrained basis, both solutions coincide at the landmark locations as shown in the figure. This example shows that the landmark specification dictates the large-scale matching between the two deer; finer matching between landmarks is found using Algorithm 2.

Next, we compare our solution to that produced by PDP. The left panel of Figure 3 shows the two optimal reparameterizations (LCGD in red, PDP in blue). The two reparameterizations appear to be very similar for $t \in [0, 0.5]$. However, as another form of comparison, one can examine their derivatives as functions of the parameter value; recall that we want the solution to be a diffeomorphism. The right panel of Figure 3 shows the corresponding two derivatives. The PDP solution is generally not smooth due to its piecewise

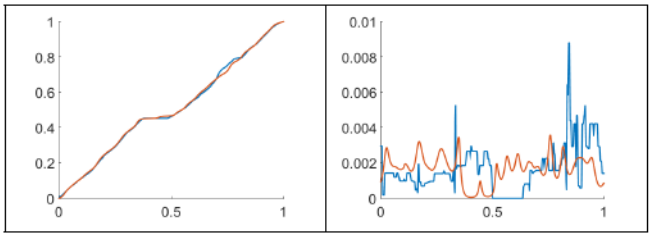

Fig. 3. Left: Optimal reparameterizations. Right: Derivative of optimal reparameterizations. LCGD=red; PDP=blue.

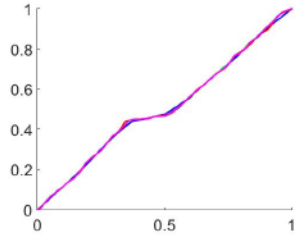

Size	Distance
5	0.6328
20	0.5876
35	0.5719
50	0.5528
65	0.5694

Fig. 4. Left: LCGD solutions using 5 (blue), 20 (red), 35 (black), 50 (green), 65 (magenta) basis elements. Right: Geodesic distance between deer using different basis sizes.

construction and other aspects of DP. Thus, this solution is not a valid element of Γ_0. On the other hand, the solution obtained using the LCGD method has a smoothly varying derivative. This helps illustrate that, while the piecewise solution appears smooth when viewed globally, zooming-in shows its local non-smoothness. In contrast, our solution is guaranteed to provide a valid minimizer in the subgroup Γ_0. Computing time for LCGD (14.24 seconds on a personal computer) is slightly longer than PDP (0.13 seconds;), but still reasonable.

The implications of using PDP versus the proposed LCGD method extend to geodesic path and distance calculations. These two quantities are heavily dependent on the registration of points across shapes, which is impacted by differences (even minuscule ones) in the reparameterizations. Figure 5 shows the resulting shape distances along with the corresponding geodesic paths. The LCGD method yields a slightly higher distance due to the smoothness requirement imposed by Γ_0, which is not enforced by PDP. The geodesic paths are also impacted; while not drastic, there are noticeable differences in how the antlers are deformed, as well as the gap between the two front legs (where the occlusion occurs). The antlers deform differently due to the difference between the two reparamterization solutions for high parameter values, which correspond to the location near the antlers.

We also study the impact of the basis approximation on the solution obtained from the proposed LCGD algorithm; truncation of the basis results in certain fine details being ignored, but one expects that there is some value for which the perceived improvement in the solution is marginal. For instance, we used 35 landmark-constrained basis elements in this example as an approximation to the full basis for $T_{\gamma_{id}}(\Gamma_0)$. To understand the effect of truncating the basis, we computed the LCGD solution for four other basis sizes: two smaller and two larger. The solutions for these five basis approximations are shown in Figure 4. In this example, the most prominent difference occurs near the end of the reparameterizations. It is clear that five basis elements are not enough to capture any of the details captured by the larger bases (this solution is substantially different from the others), and even 20 may not be enough. However, it appears that using 35, 50, or 65 basis elements all seem to result in very similar solutions. This is a direct consequence of the basis construction, where some elements oscillate much more, allowing one to capture finer details when matching shapes. These similarities also carry through when calculating the geodesic distances between shapes. When at least 35 basis elements are used, distances are very similar, as shown in the table in Figure 4.

Camel Example. Figure 6 shows a comparison of two camels, where the gaps between both the front and back legs are occluded. There are also differences in the structure of the humps and the tail. Because of the occlusion, placing landmarks is beneficial to our analysis. We used eight landmarks, with a total of $N = 800$ points. We obtain similar solutions using PDP and LCGD, but again, the proposed solution is guaranteed to be an element of Γ_0. Figure 5 also shows the distances and geodesic paths for the two solutions. Again, differences are visible in the deformation of the front and back legs. In this case, the LCGD approach significantly outperforms the PDP algorithm in terms of the distance between the registered shapes. LCGD converges in 55.85 seconds, as compared to 0.19 seconds for PDP.

Comparison on Larger Datasets. We justify the use of LCGD (due to its search for a smooth solution) by evaluating two quantities of the obtained PDP and LCGD solutions for two classes of shapes within the MPEG-7 dataset: 20 bones and 20 horses. Each bone was manually annotated with four landmarks, while each horse was marked with six landmarks. The PDP and LCGD solutions were obtained for all possible comparisons within each shape class (i.e., 190 registrations per shape class). To evaluate smoothness, we use two measures: (1) a local measure: average maximum value of the derivative of the reparameterizations over all comparisons; and (2) a global measure: average squared norm of the reparameterization derivatives over all comparisons. The results are provided in Table 1. As expected, we observe a smaller maximum derivative on average using LCGD since the algorithm searches for smooth solutions directly in Γ_0. The global measure also results in lower values for both shape classes using LCGD. These results clearly demonstrate the benefits of LCGD over PDP.

IV. Summary and Future Work

The piecewise dynamic programming approach for matching landmark-constrained shapes approximates the solution well in many cases. However, the piecewise nature of the resulting reparameterization violates the smoothness requirements of Γ_0, the subgroup of landmark-constrained diffeomorphisms. In this paper, we presented a gradient descent algo-

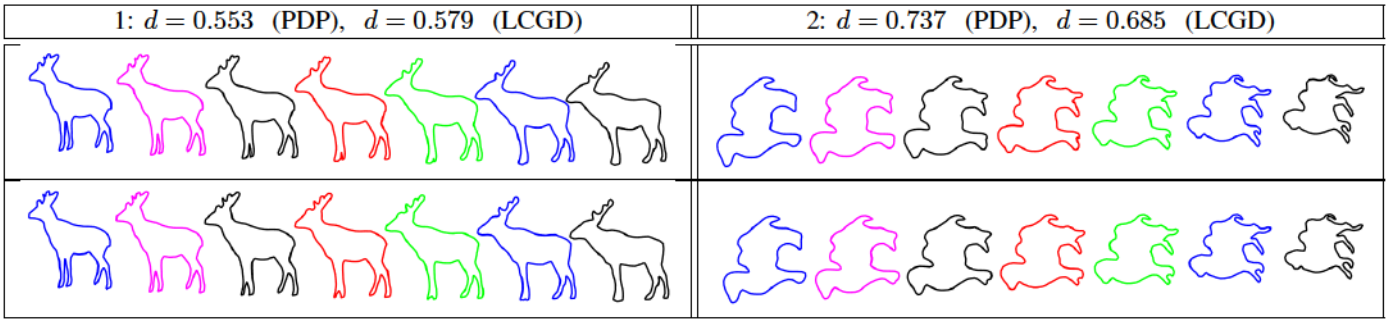

Fig. 5. Geodesic path and distance of (1) deer and (2) camels for PDP (top) and LCGD (bottom).

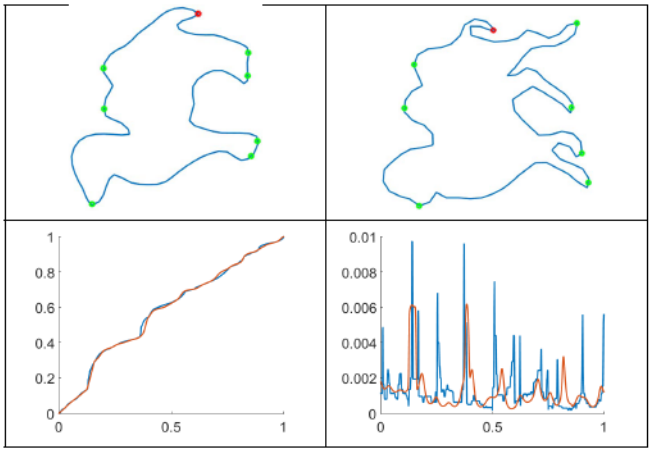

Fig. 6. Top: Camel outlines with marked landmarks (green and red). Bottom: Optimal reparameterizations (left) and their derivatives (right). LCGD=red; PDP=blue.

	Max Derivative		Squared Norm of Derivative	
Class	PDP	LCGD	PDP	LCGD
Bones	0.0102	**0.0056**	1.04×10^{-5}	$\mathbf{6.83 \times 10^{-6}}$
Horses	0.0116	**0.0096**	5.17×10^{-6}	$\mathbf{4.93 \times 10^{-6}}$

Table 1. Comparison of smoothness of registrations produced by PDP and LCGD. Best performance is highlighted in bold.

rithm, which provides an alternative approach to finding the optimal matching between landmark-constrained shapes while satisfying this smoothness requirement. We demonstrate, using several examples, that the proposed solution is often different from the PDP solution (suggesting that the PDP solution may not be an element of Γ_0). We also provide multiple quantitative and qualitative assessments of the proposed algorithm.

In many applications, having a sufficiently smooth registration is important. Consider handwriting recognition, where the task may be to detect forgeries from a batch of similar signatures. The parameterization of a signature dictates the speed at which the writer signs the name. To compare such signatures, one may want to take into account such speed variations. Methods for analyzing parameterization functions often require a certain level of smoothness [8].

A possible extension of the proposed method is to use the Riemannian geometry of Γ_0 to speed-up convergence of the proposed optimization algorithms. In the unconstrained case (where no landmarks are present), a solution to this problem can be developed under a different representation of reparameterizations, which can be shown to lie on the unit Hilbert sphere. Thus, any updates to the current estimate of the reparameterization can be found using fast geometric tools.

REFERENCES

[1] B. Avants and J. Gee. Formulation and evaluation of variational curve matching with prior constraints. In *Proceedings of the International Workshop on Biomedical Image Registration*, 2003.

[2] M. Beg, M. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, 2005.

[3] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings of the KDD Workshop*, 1994.

[4] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis, with Applications in R. Second Edition.* John Wiley and Sons, Chichester, 2016.

[5] S. C. Joshi and M. I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8):1357–1370, 2000.

[6] S. H. Joshi, E. Klassen, A. Srivastava, and I. H. Jermyn. A novel representation for Riemannian analysis of elastic curves in $\mathbb{R}^n$. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7, 2007.

[7] D. G. Kendall. Shape manifolds, Procrustean metrics, and complex projective shapes. *Bulletin of London Mathematical Society*, 16:81–121, 1984.

[8] S. Kurtek and A. Srivastava. Handwritten text segmentation using elastic shape analysis. In *Proceedings of the International Conference on Pattern Recognition*, pages 2501–2506, 2014.

[9] S. Kurtek, A. Srivastava, E. Klassen, and Z. Ding. Statistical modeling of curves using shapes and related features. *Journal of the American Statistical Association*, 107(499):1152–1165, 2012.

[10] S. Kurtek, A. Srivastava, E. Klassen, and H. Laga. Landmark-guided elastic shape analysis of spherically-parameterized surfaces. *Computer Graphics Forum (Proceedings of Eurographics)*, 32(2):429–438, 2013.

[11] R. S. Palais. Morse theory on Hilbert manifolds. *Topology*, 2(4):299 – 340, 1963.

[12] D. T. Robinson. *Functional Data Analysis and Partial Shape Matching in the Square Root Velocity Framework*. PhD thesis, Florida State University, 2012.

[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.

[14] A. Srivastava and E. P. Klassen. *Functional and Shape Data Analysis.* Springer, New York, 2016.

[15] J. Strait, S. Kurtek, E. Bartha, and S. MacEachern. Landmark-constrained elastic shape analysis of planar curves. *Journal of the American Statistical Association*, DOI: 10.1080/01621459.2016.1236726, 2016.

[16] L. Younes. Computable elastic distance between shapes. *SIAM Journal of Applied Mathematics*, 58(2):565–586, 1998.