Bicycle Chain Shape Models

Stefan Sommer [1,2], Aditya Tatu [2], Chen Chen [2], Dan R. Jørgensen [2,3], Marleen de Bruijne [2,4], Marco Loog [2,5], Mads Nielsen [2,3], François Lauze [2]

Abstract

In this paper we introduce landmark-based preshapes which allow mixing of anatomical landmarks and pseudo-landmarks, constraining consecutive pseudo-landmarks to satisfy planar equidistance relations. This defines naturally a structure of Riemannian manifold on these preshapes, with a natural action of the group of planar rotations. Orbits define the shapes. We develop a Geodesic Generalized Procrustes Analysis procedure for a sample set on such a preshape spaces and use it to compute Principal Geodesic Analysis. We demonstrate it on an elementary synthetic example as well on a dataset of manually annotated vertebra shapes from X-ray. We re-landmark them consistently and show that PGA captures the variability of the dataset better than its linear counterpart, PCA.

1. Introduction

There is a wide literature on shape representation and shape analysis in Computer Vision and Medical Imaging as shape understanding is one of the most fundamental task in Image Analysis. A 2-dimensional shape is generally defined as an equivalence class of smooth 1-dimensional submanifolds of $\mathbb{R}^2$ modulo similarity [13]. Computational representations, ranging from the simplest to the most sophisticated, have been suggested in the past, e.g. point set distributions [9, 1], linear point distribution models (PDM) [4], parametric representations via B-splines, levelset representations [16], and their adaptation, as for example, specific shape constraints, soft priors, *etc*..., for an ever growing amount of tasks.

Manual annotations of anatomical structures in medical images, such as X-rays, Ultra Sound, are routinely performed by radiologists and other experts in many clinical studies, resulting in the encoding of shapes as point set distributions. Point set distributions for shape representation and analysis are of tremendous importance in Medical Imaging. Deriving such distributions presupposes consistent annotations, which is not always the case: the following figure shows two annotated vertebra shapes from X-ray images, during a clinical study on vertebra fractures, the first vertebra is annotated with 31 points, the second with 32. Moreover the number of points between corner landmarks (the circular ones) do not match for corresponding pairs. This is caused by the absence

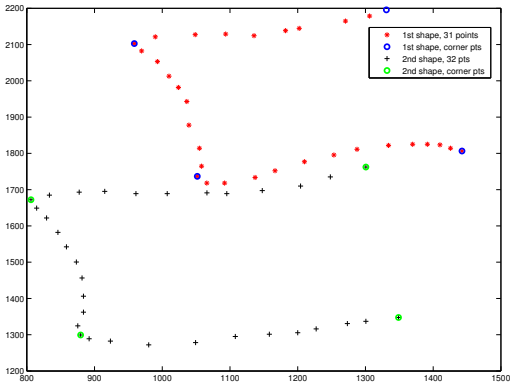

Figure 1. Two annotated vertebrae from a clinical study. The number of annotation points differ.

of clear ground truth landmarks along the endplates of the vertebrae. In order to tackle this somehow common situation, a resampling is necessary; pseudo-landmarks should be placed such that the resulting model is more compact, no additional variation caused by points sliding along the outline should be modelled. Some recent approaches for curves and surfaces were proposed for instance by Davies *et al*. [5] using minimum description length to solve this problem, while, for surfaces, Cates *et al*. used an entropy based particle system approach in [3].

For curves, which are the objects of interest in

[1]Corresponding author email: sommer@diku.dk
[2]Department of Computer Science, University of Copenhagen, Denmark
[3]Nordic Bioscience A/S, Herlev, Denmark
[4]Erasmus MC, Rotterdam, The Netherlands
[5]ICT Group, Delft University of Technology, The Netherlands

978-1-4244-3993-5/09/$25.00 ©2009 IEEE

this paper, a way to do it is to first impose a fixed number of pseudo-landmarks between landmarks, regularly distributed along the outline between the landmarks. This regularity often takes the form of an equidistance constraint for pseudo-landmarks situated between consecutive landmarks. This can be formulated as setting the variance of the distribution of of planar distances (or square-distances) between consecutive pseudo-landmarks to 0. In a figurative way, a segment between two consecutive landmarks is similar to a segment of a **bicycle chain**, for the links that constitute a bicycle chain have the same length! This has the nice property of minimizing the variability due to the annotation process. But once this resampling has been performed, forgetting this variance constraint induces apparent extra variability which may be difficult to handle due to the non linearity of the constraint. This is illustrated in Figure 2 where the Euclidean mean of the upper and lower curves does not have equidistant pseudo-landmarks introducing extra variability on the horizontal placement of the pseudo-landmarks. We

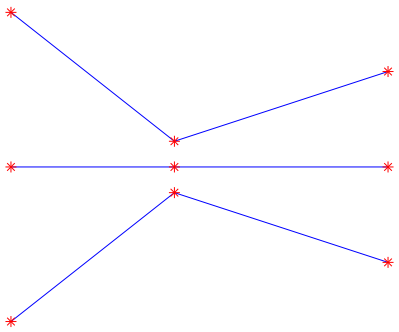

Figure 2. Two 3-point curves and the Euclidean mean.

propose to handle this situation by introducing the constraint explicitly in the descriptions of our preshape spaces. This null-variance can be reformulated as a series of simple quadratic constraints on the pseudo-landmarks and will, for shapes determined by n points in $\mathbb{R}^d$, define implicitly a submanifold of the point set spaces $(\mathbb{R}^d)^n$. Endowed with the metric that comes from the standard Euclidean structure of $(\mathbb{R}^d)^n$, it becomes a Riemannian manifold. In order to obtain point distributions models, Generalized Procrustes Analysis (GPA) [8] should be performed with the induced metric, leading to what we will call Geodesic Generalized Procrustes Analysis (GGPA), while Principal Component Analysis should be replaced by Principal Geodesic Analysis (PGA) [7] in order to take into account the curved structure of the manifold. In the rest of this paper, we will focus to point set configurations in $\mathbb{R}^2$. This will simplify the presentation. Extension to 3D curves can be carried out easily.

So as to be able to compute GGPA and PGA, we need tools for computing Riemannian exponential map, geodesics, and log map on implicitly defined submanifolds. By extending computations of exponential map to provide not only geodesic, but corresponding moving frames, we propose a shooting method for computing Log maps. When it fails, we replace it by a path straightening algorithm based on local properties of geodesics.

This paper is organized as follows. In the next section we introduce the preshape manifolds that we use as well as the geometric tools needed for our statistical analysis: Geodesic Generalized Procrustes analysis and Principal Geodesic Analysis. Exponential and Log maps are discussed in Section 3. In Section 4 we present experiments; the first one on the 3-points toy example and the second on a data set of vertebra coming from a clinical study on vertebra fractures. Finally we conclude in Section 5.

2. Preshape manifolds

In point based models, a typical object consists of q landmark points and n_k, $k = 1, ..., q-1$ ($k = 1, ..., q$, for closed configurations) pseudo landmarks between consecutive landmark points. A segment of this object consists of $n_k + 2$ points, n_k pseudo-landmarks $P_i, i = 2, ..., n_k + 1$ between 2 landmark endpoints P_1, P_{n_k+2}. The objects we consider consist of such configurations with equal (squared) Euclidean distance between the neighboring points in each of the segments. This characteristic distance will generally vary from segment to segment and objects to objects, even when the sequence of numbers $(q, n_1, \ldots, n_{q-1})$ is fixed. We start by describing constraints on segments.

2.1. n-Links Bicycle Chain Manifolds

Here onwards we work on one segment with $n_k = n-2$ pseudo-landmarks between 2 landmark endpoints. Then the equidistant constraint can be written as a simple quadratic constraint $F : \mathbb{R}^{2n} \to \mathbb{R}^{n-2}$ given as

$$F_i(P_1, ..., P_n) = d_{i+2,i+1} - d_{i+1,i}, \quad i = 1, .., n-2 \quad (1)$$

where $d_{i,j} = (x_i - x_j)^2 + (y_i - y_j)^2$ is the squared euclidean distance between points P_i and P_j, The configuration space is the subspace of $\mathbb{R}^{2n}$ given by $A_n = F^{-1}(0) \backslash \Delta$, where Δ is the "diagonal" $\Delta = (P, \ldots, P) \in (\mathbb{R}^2)^n$ consisting of segments reduced to a single point, for, while $\Delta \subset F^{-1}(0)$, the rank of F breaks down exactly along Δ. This ensures that A_n is

a submanifold of $(\mathbb{R}^2)^n = \mathbb{R}^{2n}$ [2] The tangent space of A_n at a segment x is given by

$$T_x A_n = \ker(JF(x))$$

the kernel (or null space) of the Jacobian of F at point $x \in A_n \subset \mathbb{R}^{2n}$. By restricting the scalar product of $\mathbb{R}^{2n}$ to $T_x A_n$, A_n is endowed with a Riemannian Metric [2]. We may call A_n a **n-links bicycle chain segment manifold**.

More general point configurations are then built by concatenating these n-links bicycle chain segments, imposing endpoint matching which are linear constraints. When the number q of landmarks points and the numbers n_k, $k = 1, \ldots, q-$ of pseudo-landmarks points are fixed, corresponding configurations form a Riemannian submanifold of the product manifold $\prod_{i=1}^{q} A_{n_k+2}$, and this manifold has also the metric inherited from the embedding space $(\mathbb{R}^2)^N$ with $N = q + \sum_{k=1}^{q-1} n_k$.

Having a Riemannian metric, we can compute length of paths in these manifolds, define geodesic and geodesic distances [2].

2.2. Removing Translation and Scaling

In the following, we denote by $\mathcal{M}$ such a configuration manifold. To work with preshapes in the sense of [9], we need to quotient out translations and scaling from points in $\mathcal{M}$ (although in some applications, scale could be an important feature of the shape). Removing translations is as usual easy. If $\mathcal{M}'$ denotes the submanifold of $\mathcal{M}$ of configurations with centroid at the origin of $\mathbb{R}^2$ then $\mathcal{M} \simeq \mathcal{M}' \times \mathbb{R}^2$, by sending a configuration $S = (S_1, \ldots, S_n)$ to $(S - \bar{S}, \bar{S}$ where $\bar{S} = \frac{1}{n}\sum_{i=1}^{n} S_i$ is the centroid of S. This decomposes $\mathcal{M}$ into two *orthogonal* factors, which imply that a geodesic path in $\mathcal{M}$ between centered objects in $\mathcal{M}'$ will be in fact a geodesic path in $\mathcal{M}'$. From now on we therefore assume that all our configurations have centroid at $0 \in \mathbb{R}^2$. Following [9], we remove scale by imposing $\|S\|^2 = \sum_{i=1}^{n} \|S_i\|^2 = 1$, i.e by intersecting $\mathcal{M}'$ with the unit sphere of the embedding space. This defines a new submanifold $\mathcal{S}$ of $\mathcal{M}'$, and $\mathcal{S}$ is our preshape manifold.

2.3. Geodesic Generalized Procrustes Analysis

Given a sample set $(S_i)_{i=1\ldots n} \in \mathcal{S}$, our GPA follows [9], but is performed on $\mathcal{S}$. It attempts to compute a set of planar rotations $R_{\bar{\theta}_i}$, $i = 1, \ldots, n$ and a preshape $\bar{\mu} \in \mathcal{S}$ minimizing the misalignment criterion

$$E(\theta_1, \ldots, \theta_n, \mu) = \sum_{i=1}^{n} d(R_{\theta_i} S_i, \mu)^2. \qquad (2)$$

where d is the geodesic distance in $\mathcal{S}$. This will result in an aligned preshape sample $(\bar{S}_i := R_{\bar{\theta}_i} S_i)_{i=1\ldots n}$, $\bar{\mu}$ being the Fréchet mean ([10]) of this sample and the distances $d(\bar{S}_i, \bar{\mu})$ should represent the true *shape* distances to this mean.

The minimization procedure for (2) is sketched in Algorithm 1. We describe briefly the loop steps. A first guess for the rotations is computed by standard Euclidean rigid registration [8] providing candidate rotation angles for each shape. Then we search for the rotations angles that minimize the true geodesic distances in a neighborhood of the previously obtained angles. The Fréchet mean is then computed by adapting the procedure described in [7] to our case. In Figure 3 the need for the minimization search after the initial Euclidean registration is illustrated by showing a base preshape, and rotation of a second preshape with respect to Euclidean and submanifold distances.

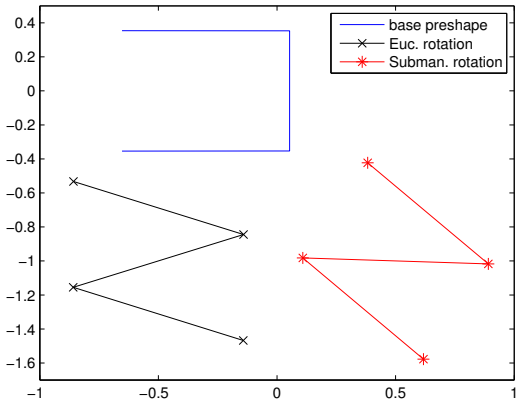

Figure 3. A base preshape and Euclidean and manifold registration.

Algorithm 1 Calculate the mean shape μ' and the aligned shapes S_i''

Require: $S_i \in \mathcal{S}$, $i = 1, ..., m$
  $\mu' = S_1$ {initial guess}
  repeat
    Set $\mu = \mu'$. $S_i' = S_i$ shapes aligned to μ using Euclidean distances.
    $S_i'' = S_i'$ shapes aligned to μ using *geodesic* distances.
    $\mu' =$ Fréchet mean of (S_i'')
  until $d(\mu, \mu') <$ Threshold.
  Output: Mean μ', aligned shapes S_i''.

2.4. Principle Geodesic Analysis (PGA)

PGA is a generalization of Principal Component Analysis (PCA) to nonlinear manifolds [7]. We seek to compute a minimum number of tangent vectors at the mean, which generate geodesics that represent as much variability in the data on the manifold as possible. Thus PGA is PCA done on the tangent space of the mean. Unfolding the manifold to this tangent space is performed by the Riemannian Log map. The algorithm can be summarized as:

- Given m preshapes in $\mathcal{S}$ and the mean preshape μ, compute $v_i = Log_\mu(S_i)$, $i = 1, ..., m$, the tangent vectors for each preshape in the tangent space at the mean.
- Compute the covariance matrix $C = \frac{1}{m}\sum_{i=1}^{m} v_i v_i^T$
- Compute the eigenvectors and eigenvalues (e_i, λ_i) of C.

The geodesic paths corresponding to the tangent vectors $e_i \in T_\mu\mathcal{S}$ are the principle geodesic components.

Computing the Fréchet mean and PGA use Exponential map, Log map and geodesics on implicitly defined Riemannian manifolds. They are described in the next section

3. Geodesics on the manifold; the Exp- and Log-map

Geodesics are fundamental to the theory of Riemannian manifolds ([2]). They are closely related to the Exponential map $\mathrm{Exp} : TM \to M$ in the sense that a geodesic γ trough the point p with initial velocity vector v is given by the curve

$$\gamma(t) = \mathrm{Exp}_p tv \ .$$

The map Exp_p is invertible in a sufficiently small neighborhood of 0 in T_pM. When U is such a neighborhood we denote by $\mathrm{Log}_p : \mathrm{Exp}_p(U) \to U$ the inverse of Exp_p.

The distance between two elements of the manifold is given by

$$d_M(p,q) = \inf\left\{l(c)\middle| c \text{ is a curve joining } p \text{ and } q\right\} \ .$$

Here $l(c)$ denotes the length of the curve c. Since geodesics are critical points of the length functional, it is in the case of a complete manifold M sufficient to consider geodesics when computing the distance. Therefore, if p and q are sufficiently close so that only one geodesic joins them,

$$d_M(p,q) = \|\mathrm{Log}_p q\| \ . \tag{3}$$

In general we cannot be sure that a given geodesic joining p and q is length minimizing. In such cases, we define $\mathrm{Log}_p q$ to be the initial direction of *some* geodesic joining p and q and use (3) as a guess on the distance.

Computing Exp_p amounts to solving an initial value ODE problem. This can be done neatly numerically, confer [6]. Computing $\mathrm{Log}_p q$ is substantially harder. We make use of a shooting method ([12], [14]) for computation of Log_p for input values close to p, and a path-straightening method for non-local input.

3.1. Shooting method

A shooting method iteratively improves an initial guess by repeatedly computing a residue or error correction, and updates the initial guess using that. Based on the fact that Log_p is the inverse of Exp_p, our basic algorithm is presented in Algorithm 2. The ability to compute length and direction in Euclidean space and the implicit representation of $\mathcal{S}$ as a submanifold of Euclidean space enables us to compute both the initial guess, update v, and compute the Euclidean error of our guess. When q is close to p these estimates work well and improve the situation in [14] where the embedding space approximations are not at hand and e.g. the update of v therefore is based on numerical estimates of the gradient of a cost functional.

Algorithm 2 Calculate $v = \mathrm{Log}_p q$ on $\mathcal{S}$ by shooting

Require: $p, q \in \mathcal{S}$
  $v \Leftarrow$ projection of $q - p$ to $T_p\mathcal{S}$ {initial guess}
  repeat
    $\tilde{q} \Leftarrow \mathrm{Exp}_p v$ {shot based on guess}
    $\tilde{r} \Leftarrow$ projection of $q - \tilde{q}$ to $T_{\tilde{q}}\mathcal{S}$ {residue at $\tilde{q}$}
    $r \Leftarrow$ par. transport of $\tilde{r}$ to $T_p\mathcal{S}$ {residue at p}
    $v \Leftarrow v + r$ {update v}
  until $\|\tilde{q} - q\|_{\mathbb{R}^{2n}}$ is sufficiently small.

We use the projection of the vector $q - p$ in embedding Euclidean space to the tangent space $T_p\mathcal{S}$ as our initial guess. In each iteration we compute $\mathrm{Exp}_p v$ and express the error by the Euclidean distance $\|q - \mathrm{Exp}_p v\|$. We update v by projecting the Euclidean residue $q - \mathrm{Exp}_p v$ onto the tangent space $T_{\mathrm{Exp}_p v}\mathcal{S}$, parallel transport the resulting vector to $T_p\mathcal{S}$ and add it to v; this procedure is the natural manifold generalization of error correction in Euclidean space.

The parallel transport is computed using a parallel frame along the curve $t \mapsto \mathrm{Exp}_p tv$. We compute the parallel frame by using the fact that parallel vector fields have zero intrinsic acceleration, introduce a Lagrange multiplier, and solve the resulting ODE. The computation of the frame can be nicely coupled with

the computation of $\text{Exp}_p v$ when using the method of [6].

The shooting method relies completely on the quality of the initial guess and updating residues. Both are determined by how well the projections on the tangent spaces approximate the paths on the manifold, or, in other words, how close to linear the manifold is; in an Euclidean manifold the shooting method converges in one iteration whether as it on a torus might not converge at all. It will though always converge locally due to the smoothness of our manifold.

An additional drawback of the shooting method is its sensitivity to numerical errors in the computation of Exp_p. This can especially be a problem if the curvature around the target element q is large, confer [11].

3.2. Path straightening

When the shooting method fails to converge due to large curvature of the manifold, we apply the path straightening method of [15]; we update an initial curve by repeatedly shooting between pairs of points on the initial curve close to each other. The closeness assures the convergence of the shooting method. In each iteration the curve is a piecewise geodesic and by repeatedly changing the points between which we shoot, the non-smooth bends of the curve are removed. Since geodesics are critical points of the length functional, we stop the process when we get no signification reduction of length on each iteration.

Path straightening requires an initial path. In practice we get this path by shooting until we detect non-convergence of the shooting method. We then restart the shooting method with the best guess from the previous run as our new starting point. In practice we always obtain convergence of the shooting method in the second run. Now concatenating the geodesics obtained from the two runs gives a piecewise geodesic connecting the points which can serve as input to the path straightening algorithm. In case this method fails, we explicitly make an initial path.

As noted in [15] we may need to extract a subsequence in order for the path straightening algorithm to convergence to a geodesic. In practice we do not experience such situations, and we accept the possibility of this happening in the same way as we accept that geodesics might not be length minimizing.

4. Experiments

We present two examples illustrating the effect of our manifold setting. We start by discussing the dimensionality reduction gained in a small 3-point example and then progress to study a dataset of vertebrae shapes.

4.1. Illustrative example

In Figure 4 we see three 3-point preshapes with equidistant points. They are all normalized and hence reside on the manifold $\mathcal{S}$. The middle preshape is the Fréchet mean of the upper and lower preshapes, and hence the mean of all three preshapes.

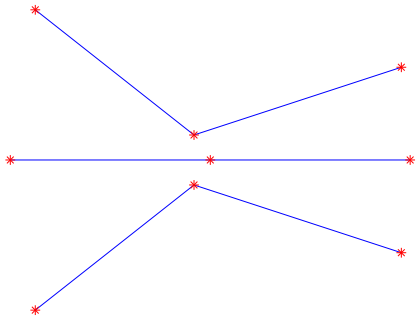

Figure 4. Three 3-point preshapes on the manifold.

The manifold $\mathcal{S}$ has two dimensions. Doing a Principal Geodesic Analysis on the set of the three preshapes, we get one mode of variation. The geodesic corresponding to this mode connects the three preshapes as illustrated in Figure 5. Note that in the figure the preshapes have been placed in the plot as to have zero mean.

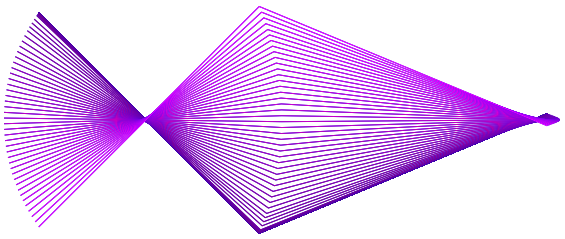

Figure 5. The geodesic corresponding to the only mode of variation obtained from PGA.

Now suppose we disregard our manifold notion and attempt to do Euclidean Principal Component Analysis in the embedding Euclidean space. The Euclidean mean of the three preshapes will again be a straight line, but in this situation the points on the mean will not be equally spaced and hence the mean will not be in $\mathcal{S}$. When computing the PCA we get two modes of variation; one mode representing vertical motion as illustrated in Figure 6, and one mode representing horizontal motion. The latter mode arises from the placement of the points on the straight line mean and is

thus irrelevant. Therefore, in this example the PCA captures only 97.5 percent of the variation in a mode giving relevant information. This contrasts that PGA captures all variation.

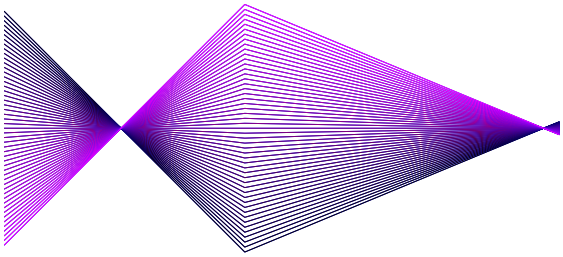

Figure 6. One of two modes of variation for PCA.

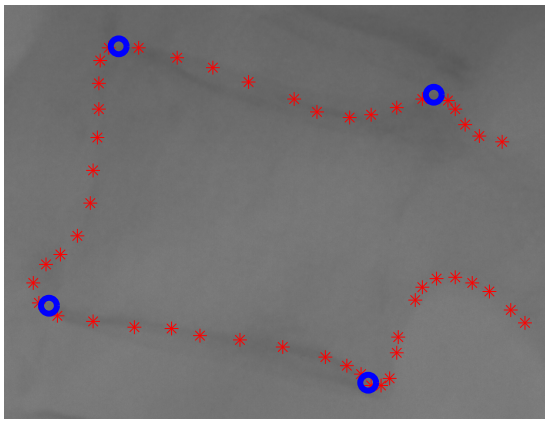

Figure 7. Manually annotated vertebra.

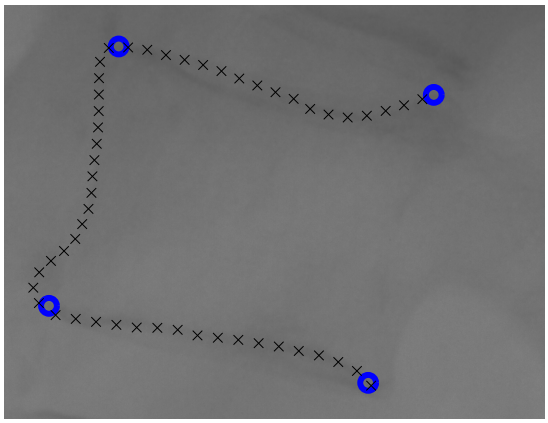

Figure 8. Result of applying redistancing procedure.

4.2. Vertebra shapes

Our dataset of vertebrae consists of 304 manually annotated vertebra shapes on lateral X-rays. For each vertebra, outlines have been manually drawn by choosing points along the contours, assuming a simple linear interpolation between them. Corner points of the vertebra endplates are indicated but do not always match the outlines perfectly. New corner points have been defined as the points of the contour that are closest to the manually annotated corners. This divides the outline into 3 segments, the upper, left and lower ones. For resampling, we fixed the number of pseudo-landmarks per segment to 16, leading to 52 points per shape. The pseudo-landmarks positions were computed segment-wise so as to minimize a squared-distance between the original outline and the new one. Given an n-tuple $\mathcal{P} = (P_1, \ldots, P_n)$ of equidistant-spaced points, with P_1 and P_n being the *fixed* corner points of this segment, let $C_{\mathcal{P}}(t)$ be the piecewise linear curve joining them, and $C_0(t)$ the piecewise linear curve formed by joining the original annotated points for the corresponding segment. We minimize the squared-distance

$$E(\mathcal{P}) = \int (C_{\mathcal{P}} - C_o)^2 \, dt.$$

We start with a configuration $\mathcal{P}$ on the straigth line segment joining P_1 to P_n and perform gradient descent on the corresponding preshape manifold $\mathcal{S}$ using the exponential map. The result of applying the redistancing procedure to the manually annotated vertebra in Figure 7 is shown in Figure 8.

In our illustrative example it is clear that we introduce non-linearity when restricting to the manifold. In order to illustrate that we have significant curvature also in the relatively high dimensional manifold used for the vertebrae, we compute the Fréchet mean vertebra and measure an approximate distance from each vertebra to the tangent space of the mean; we let v_m denote the mean and for each vertebra v we compute $w = \mathrm{Log}_{v_m} v$. We then let x be the distance $\|v - (v_m + w)\|$ between the vertebra and an approximated projection to $T_{v_m}\mathcal{S}$, and record the relative distance $x/\|w\|$. A non-curved manifold would result in zero relative distance. We see a mean relative distance of 12 percent clearly indicating that the manifold is curved. Performing the same computation on the non-normalized manifold $\mathcal{M}'$ gives a mean relative distance of 9 percent indicating that not all curvature arises from the normalization to the unit sphere.

Figure 9 illustrates how PGA provides a more compact description than PCA. The figure shows the normalized sum of the first n eigenvalues as a function of n. It can be seen that in order to capture say 99.5 percent of the variation, we will need 25 eigenvectors

when doing PCA as opposed to only 20 eigenvectors when doing PGA.

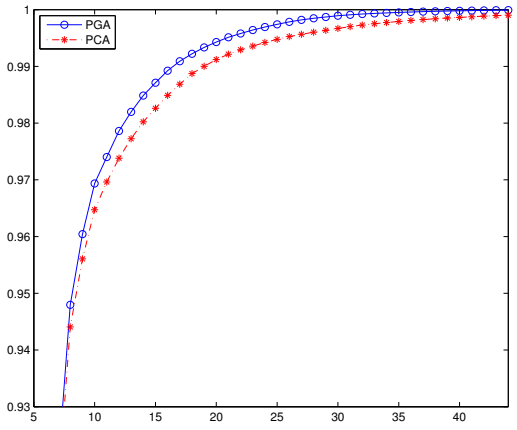

Figure 9. Accumulated spectrum of PGA and PCA.

5. Conclusion

In this paper we have introduced manifolds of preshapes built by constraining distributions of pseudo-landmarks between pairs of consecutive landmarks. This endows these preshape manifolds with a structure of Riemannian manifolds. We have developed tools for computing Exponential maps, Log maps, geodesic distances, allowing us to define a Geodesic GPA and adapt PGA to that situation. We have shown on examples that PGA captures variability better that PCA.

Although we have built our models for planar point configurations, they are clearly not restricted to this case. Other types of length and position constraints can also be used. We are also not restricted to shape manifolds. The techniques presented in this work can be used to perform statistics on other submanifolds of a linear configuration space implicitly defined by a set of smooth constraints. This is the subject of ongoing work.

References

[1] F. Bookstein. *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambrigde University Press, 1991.

[2] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry. Revised Second Edition*. Academic Press, 2003.

[3] J. Cates, P. T. Flechter, M. Styner, M. Shenton, and R. Whitaker. Shape Modeling and Analyis with Entropy-Based Particle Systems. In B. P. F. L. Nico Karssemeijer, editor, *Proceedings of the 20th International Conference on Information Processing in Medical Imaging*, volume 4584 of *LNCS*, pages 333–345, Kerkrade, The Netherlands, 2007. Springer.

[4] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[5] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. A Minimum Description Length Approach to Statistical Shape Modeling. *IEEE Transactions on Medical Imaging*, 21(5):525–538, may 2002.

[6] J. Dedieu and D. Nowicki. Symplectic methods for the approximation of the exponential map and the newton iteration on riemannian submanifolds. *Journal of Complexity*, 21(4):487–501, Aug. 2005.

[7] P. T. Fletcher, S. Joshi, C. Lu, and S. M. Pizer. Principal Geodesic Analysis for the Study of Nonlinear Statistics of Shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005, Aug 2004.

[8] C. Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society B*, 53(2):285–339, 1991.

[9] D. Kendall. Shape manifolds, procrustean metrics and complex projective spaces. *Bulletin of London Mathematical Society*, 16:81–121, 1984.

[10] D. G. Kendall, D. Barden, T. K. Carne, and H. Le. *Shape and Shape Theory*. Wiley and Sons, 1999.

[11] E. Klassen and A. Srivastava. *Geodesics Between 3D Closed Curves Using Path-Straightening*, volume 3951 of *Lecture Notes in Computer Science*, pages 95–106. Springer, 2006.

[12] E. Klassen, A. Srivastava, W. Mio, and S. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:372—383, 2004.

[13] P. W. Michor and D. Mumford. An overview of the riemannian metrics on spaces of curves using the hamiltonian approach. *math/0605009*, Apr. 2006. Applied and Computational Harmonic Analysis 23 (2007), 74-113.

[14] W. Mio and A. Srivastava. Elastic-string models for representation and analysis of planar shapes. *Proceedings of the IEEE Computer Society International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2:10—15, 2004.

[15] L. Noakes. A global algorithm for geodesics. *Journal of the Australian Mathematical Society*, 64:37–50, 1998.

[16] J. Sethian. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Sciences*. Cambridge Monograph on Applied and Computational Mathematics. Cambridge University Press, 1999.