

An Investigation of Projective Shape Space

John T. Kent

1 Introduction

Consider a configuration $X_0(k \times m)$ of k points or landmarks in m -dimensional space. By identifying configurations which are related to one another by a certain group action, we obtain the concept of a “shape” as an equivalence class of configurations. The collection of equivalence classes forms a “shape space”. Here are several important examples.

1. *Similarity shape space.* The similarity shape of X_0 can be described as the equivalence class of configurations

$$[X_0]_{\text{SS}} = \{\beta X_0 R + 1_k \gamma^T : \beta > 0, R \in SO(m), \gamma \in R^m\},$$

under similarity transformations; β is a scaling parameter, R represents an $m \times m$ rotation matrix, and γ represents a translation parameter. See, e.g., [2];

2. *Reflection similarity shape space.* As above, but now suppose R is an orthogonal matrix (so reflections are allowed). The reflection similarity shape of X_0 can be described as the equivalence class of configurations

$$[X_0]_{\text{RSS}} = \{\beta X_0 R + 1_k \gamma^T : \beta > 0, R \in O(m), \gamma \in R^m\}$$

(e.g. [3]);

J.T. Kent (✉)

Department of Statistics, University of Leeds, Leeds LS2 9JT, UK
e-mail: j.t.kent@leeds.ac.uk

3. *Affine shape space.* Replacing βR by a general nonsingular matrix yields an affine shape as the equivalence class of configurations

$$[X_0]_{AS} = \{X_0 A + 1_k \gamma^T : A(m \times m) \text{ nonsingular, } \gamma \in R^m\};$$

4. *Projective shape space.* Let $P\mathcal{S}(k, m)$ denote the projective shape space of k landmarks in m dimensions. To describe this space, it is helpful to switch to homogeneous coordinates. Introduce an augmented matrix

$$X = [X_0 \ 1_k],$$

where 1_k is a k -vector of ones. Then X is a $k \times p$ matrix, where throughout the paper we set $p = m + 1$. If X is written in terms of its rows as

$$X = \begin{bmatrix} x_1^T \\ \vdots \\ x_k^T \end{bmatrix},$$

then from the point of view of homogeneous coordinates each row x_i^T of X is well-defined only up to a scalar multiple. Then the projective shape of X is defined as the equivalence class of matrices (in homogeneous coordinates)

$$[X]_{PS} = \{DXB^T : D(k \times k) \text{ diagonal nonsingular, } B(p \times p) \text{ nonsingular}\}.$$

Projective geometry is important in computer vision for identifying features in images which are invariant under the choice of camera view ([4, 6]). The matrix B holds information about the location of the focal point of the camera and its orientation. The matrix D is present because in a camera image of a point, it is not possible to determine how far away that point is in the real world.

Of course, as the group of transformations get larger, the number of distinct equivalence classes gets smaller. Unfortunately, working with equivalence classes is rather awkward from statistical point of view. Therefore various superimposition methods have been developed to facilitate quantitative comparisons between shapes. The most successful class of superimposition methods goes under the name of Procrustes analysis.

2 Review of Procrustes Methods

Consider a transformation group \mathcal{G} , with elements denoted by g , acting on configurations X_0 by taking X_0 to $g(X_0)$. Suppose that \mathcal{G} can be split into a product of three subgroups, $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ in such a way that each $g \in \mathcal{G}$ can be decomposed

(in at least one way) as $g(X_0) = g_3(g_2(g_1(X_0)))$. In any particular application, one or more of the subgroups might be trivial. Write $g = (g_1, g_2, g_3)$ to represent this decomposition. Then, as discussed in [8], the Procrustes approach to shape analysis involves several steps.

1. *Standardization.* Remove the transformation parameters in \mathcal{G}_1 by standardization. For example, if \mathcal{G}_1 denotes the location-scale group, so $g_1(X_0) = \beta X_0 + 1_k \gamma^T$ for some $\beta > 0$ and $\gamma \in R^m$, it is common to choose X_0 from the equivalence class so that it is centered and scaled,

$$X_0^T 1_k = 0_m, \quad \text{tr}(X_0^T X_0) = 1; \quad (1)$$

2. *Embedding.* Embed the standardized shape into some Euclidean space in such a way as to remove the parameters in \mathcal{G}_2 . That is, consider a mapping $\phi(X_0) = T$, say, where ϕ has the property that $\phi(X_0) = \phi(g_2(X_0))$ for all g_2 and all standardized configurations X_0 ;
3. *Optimization.* Define a (partial) Procrustes distance between the shapes of the configurations $X_0^{(1)}$ and $X_0^{(2)}$ by minimizing the Euclidean distance between the embedded objects $T^{(1)} = \phi(X_0^{(1)})$ and $T^{(2)} = \phi(g_3(X_0^{(2)}))$ over the remaining transformation parameters in \mathcal{G}_3 ,

$$d_{\text{pp}}^2(T^{(1)}, T^{(2)}) = \min_{g_3} \text{tr} \left\{ (T^{(1)} - T^{(2)})^T (T^{(1)} - T^{(2)}) \right\}; \quad (2)$$

4. *Metric comparisons.* The Procrustes distance (2) can be used directly to compare different shapes. Alternatively, as a slight variant, its infinitesimal version can be used to define a Riemannian metric on shape space and Riemannian distance can be used. In particular, each of the shape spaces under consideration can be viewed as a Riemannian manifold, other than perhaps at some singular points.

For this construction to be useful, two properties must be checked:

- *Symmetry:* $d_{\text{pp}}^2(T^{(1)}, T^{(2)}) = d_{\text{pp}}^2(T^{(2)}, T^{(1)})$;
- *Identifiability:* $d_{\text{pp}}^2(T^{(1)}, T^{(2)}) > 0$ for distinct shapes; that is, d_{pp} is a distance and not just a semi-distance.

Here are details for the three shape spaces described above.

For similarity shape, the standardization step involves centering ($X_0^T 1_k = 0_m$) and scaling ($\text{tr}(X_0^T X_0) = 1$) as in (1). The embedding step is trivial, $\phi(X_0) = X_0$. Only the rotation parameter remains at the optimization stage. For the special case $k = 3$, $m = 2$, it is well-known that the resulting shape space (shapes of triangles) can be identified with the usual sphere of radius $1/2$ in R^3 . A variant of partial Procrustes distance, known as full Procrustes distance, corresponds to chordal distance on the sphere, and Riemannian distance corresponds to great circle distance. For the purposes of this paper, these are called “Level 1” metrics.

For reflection similarity shape, there are two possible Procrustes approaches. In the first, essentially the same steps as in the previous paragraph can be applied,

with g_3 now denoting an orthogonal transformation (including reflection as well as rotation) and again leading to a Level 1 metric. However, as an alternative approach, it is more elegant at the embedding step to set $T = \phi(X_0) = X_0 X_0^T$ in terms of a configuration X_0 standardized as in (1), so that T is a $k \times k$ positive semi-definite symmetric matrix from which X_0 can be recovered up to an orthogonal transformation on the right. Hence the optimization step is not needed here. Euclidean distance between the embedded configurations will be called a “Level 2” metric in this paper and has been studied by Dryden et al. [3]. The Level 1 and Level 2 metrics are different from one another, even infinitesimally, especially when X_0 is singular or nearly singular.

For affine shape, the standardization step involves centering ($X_0^T 1_k = 0_m$) and orthonormalization ($X_0^T X_0 = I_m$), e.g. by Gram-Schmidt. The simplest embedding is given by $T = X_0 X_0^T$, which again removes the orthogonal transformations, so that no transformation parameters remain at the optimization step. Affine shape space can be identified with the Grassmann manifold of m -dimensional subspaces of R^{k-1} ($k-1$ rather than k to allow for the centering). Euclidean distance between the embedded configurations will be called “Grassmann Euclidean” distance and is another example of a Level 2 metric.

Affine shape space is already very familiar to statisticians from multiple linear regression analysis, where X_0 , taken to be centered for simplicity, plus a column of ones represents the design matrix. If y is a centered k -vector of responses, then the ordinary least squares fit of y on X_0 is given by $\hat{y} = X_0(X_0^T X_0)^{-1} X_0^T y$, which is unchanged if X_0 is replaced by $X_0 A$ for any nonsingular $p \times p$ matrix A . Note that X_0 and $X_0 A$ have the same column space, so that \hat{y} depends only on the span of the columns of X_0 , not on the individual columns themselves.

For projective shape, the standardization is more delicate. As shown in [8], it is possible to find a diagonal matrix D and a nonsingular matrix B such that after standardization the rows of the augmented configuration $X = [x_1, \dots, x_k]^T$ are unit vectors ($x_i^T x_i = 1$) and the columns are orthonormal up to a scale factor ($X^T X = (k/p)I_p$). This standardization is known as “Tyler standardization” after [11, 12]. After this standardization, X is unique up to multiplication on the left by a diagonal matrix of plus and minus ones, and on the right by an orthogonal matrix. A nice way to remove these remaining indeterminacies at the embedding stage is to define $k \times k$ matrices M and N with entries

$$m_{ij} = |x_i^T x_j|, \quad n_{ij} = (x_i^T x_j)^2 = m_{ij}^2. \quad (3)$$

Then there are two versions of Procrustes distance between the projective shapes: the Euclidean distance between the M matrices (a Level 2 metric), or between the N matrices (a Level 4 metric), respectively. At least for $m = 1$, that is, $p = 2$, it can be shown that these constructions are identifiable.

One way to think about Tyler standardization for general k and m is in terms of a camera image of a scene of k points in m dimensions affinely situated in an “ambient space” R^p , $p = m + 1$. Tyler standardization is equivalent to using a camera with spherical film (rather than the more conventional flat film) and choosing the focal

point chosen so that the moment of inertia of the film image is proportional to the identity matrix I_p .

It is worth commenting on the naming conventions for the different levels of metric. A Level 1 metric involves direct comparisons between standardized configurations (after optimizing over the remaining transformation parameters). A Level 2 metric involves comparisons between second order moments of a standardized configuration, such as $X_0 X_0^T$, (after optimizing over any remaining transformation parameters). Finally the Level 4 metric involves fourth order moments of the original configurations (with no remaining transformation parameters in our projective shape example in $m = 1$ dimension). There is no Level 3 metric in this naming system.

3 Singularities

Here is a brief summary of the singularities that arise in the different shape spaces. Here X_0 denotes a standardized (i.e. centered and scaled) configuration, and X denotes a Tyler standardized augmented configuration. The nature of any singularities depends on the particular shape space and on the metric used.

1. *Similarity shape space, partial or full Procrustes distance.* When $m = 1, 2, k \geq 3$, there are no singularities. All that is required is that X_0 be at least a rank 1 matrix. In particular, if $m = 2$ there is no singularity when X_0 has rank $r = 1$. Indeed for $m = 2$, similarity shape space is homogeneous in the language of differential geometry, meaning that every point in shape space looks like every other point.

However, singularities do arise in similarity shape space in dimensions $m \geq 3$ at configurations X_0 of rank $r \leq m - 2$. The simplest example is given by a set of collinear points in R^3 ; the singularity arises because the configuration is unchanged under a rotation in R^3 that leaves the axis of the line fixed. The high curvature near such points in shape space has been studied in [7];

2. *Reflection similarity shape space, Level 2 metric.* This space has more singularities. For all dimensions $m \geq 2$, there is a singularity whenever X_0 has rank $r < m$. So a set of collinear points in the plane is viewed as singular under the Level 2 metric for reflection similarity shape space, but not under the Level 1 metric for similarity shape space;
3. *Affine shape space, Grassmann Euclidean metric.* Here by definition the standardized configurations X_0 are assumed to have rank m . There are no singularities;
4. *Projective shape space, Level 1, 2 and 4 metrics.* Here the situation is more complicated. It is easiest to study for $m = 1$, where the Level 1, 2 and 4 metrics all have singularities at the same points in shape space [8].

4 Projective Shape Space $P\mathcal{S}(4, 1)$

Projective shape spaces are considerably more complicated than similarity or affine shape spaces. In this section we summarize some of the challenges which appear even in the simplest case $P\mathcal{S}(4, 1)$ of $k = 4$ collinear points in $m = 1$ dimension. Let the positions of the landmarks be given by 4 numbers, u_j , $j = 1, \dots, 4$. Then the projective shape can be described in terms of a single number known as the cross ratio, one version of which is defined by

$$\tau = \frac{(u_1 - u_2)(u_3 - u_4)}{(u_1 - u_3)(u_2 - u_4)}.$$

Each value of τ represents a different projective shape as τ ranges through the extended real line (with the limits $\pm\infty$ identified with one another).

In [8], it is emphasized that the cross ratio representation of projective shape is not suitable for metric comparisons. Instead various Procrustes distances are considered. Here is a sketch of the main results.


The M and N versions of Procrustes distance in (3) give rise to two simple geometric representations of $P\mathcal{S}(4, 1)$. Under the M representation, $P\mathcal{S}(4, 1)$ becomes a spherical equilateral triangle, most easily visualized as two great circle arcs along lines of longitude from the north pole to the equator separated by 90° , together with an arc along the equator connecting them. Each edge of this spherical triangle is an arc of length 90° .

For the N representation, $P\mathcal{S}(4, 1)$ becomes a planar equilateral triangle, with each edge of length 1, say. A plot is given by the outline of the central triangle given in Fig. 4. One edge, AB, say, corresponds to the interval $\tau \in [0, 1]$. The next edge, BC, corresponds to the interval $\tau \in [1, \infty]$, or equivalently, $(\tau - 1)/\tau \in [0, 1]$. The final edge, CA, corresponds to the interval $\tau \in [-\infty, 0]$, or equivalently, $1/(1 - \tau) \in [0, 1]$. The relevance of these nonlinear mappings of τ can be explained in terms of the effect on τ of various permutations of the labels of the landmarks.

However, the neatness of these geometrical representations hides some of the more subtle aspects of projective shape:

1. visualizing when different configurations have the same projective shape is intuitively difficult. Figure 1 illustrates several configurations for which the outer landmarks u_1 and u_4 are held fixed, but the inner two landmarks vary in such a way that the cross ratio remains fixed at $\tau = 0.3$. The human eye is not very good at recognizing that these configurations have the same cross ratio. (On the other hand the human eye is excellent at deducing depth information from stereo images!);
2. Tyler standardization provides a mathematically elegant way to standardize a configuration. However, in real world applications a film image of an underlying configuration will be observed subject to errors (either in the location of the landmarks in the ambient space R^p or in the location of the landmarks on the film image). Unfortunately the way these errors influence the distribution of projective

Fig. 1 Various configurations of four collinear points with the same cross ratio



shape depends on the “pose” of the original configuration in the ambient space in a considerably more complicated manner than is the case for similarity and affine shapes. A partial analysis is given in [8];

3. the singularities in projective shape space are well-defined mathematically, but are somewhat unexpected intuitively. In the current setting, $k = 4, m = 1$, they correspond to either a single pair coincidence (e.g. $u_1 = u_2 < u_3 < u_4$) or a double pair coincidence (e.g. $u_1 = u_2 < u_3 = u_4$). On the other hand, the singular points do not look very special when looked at in terms of the cross ratio;
4. the implications on statistical understanding of different metrics, e.g. those based on either Level 2 or Level 4 Procrustes analysis, is still not entirely clear;
5. a related issue is the difficulty in constructing useful and tractable models on projective shape space. The next sections give some initial suggestions.

5 Uniform Distributions on $P\mathcal{S}(4, 1)$

In this section we explore possible uniform distributions on $P\mathcal{S}(4, 1)$. In terms of the planar triangle representation, there are at least three general approaches. For each approach the density is the same on each side of the triangle and is symmetric about the midpoint of each side. Hence we limit attention to one side, parameterized by $0 < \tau < 1$. The form of the three densities is given as follows.

1. (Independent sampling) Take four independent points from a specified distribution on the line and compute the resulting distribution of the cross ratio. The distribution under normality was worked out by Maybank [9, 10] and the distribution under uniformity was worked out by [1]. The latter distribution is more tractable to write down here and can be expressed as

$$f_{\text{Ind}}(\tau) = f_0(\tau) + f_0(1 - \tau), \quad 0 < \tau < 1,$$

where

$$f_0(\tau) = \{(\tau + 1) \log \tau + 2(1 - \tau)\}/(\tau - 1)^3;$$

- (Level 2 metric) Consider a uniform distribution on the spherical triangle representation of $P\mathcal{S}(4, 1)$. Thus on each edge, the angle $\delta \in (0, \pi/2)$ is uniformly distributed. After changing the variable from δ to $\tau = \sin^2 \delta$ on the edge for $0 < \tau < 1$, the density becomes

$$f_{L2}(\tau) = 1/\{\pi\sqrt{\tau - \tau^2}\}, \quad 0 < \tau < 1;$$

- (Level 4 metric) Consider a uniform distribution on the planar triangle representation of $P\mathcal{S}(4, 1)$, so

$$f_{L4}(\tau) = 1, \quad 0 < \tau < 1.$$

In each case the density has been scaled to integrate to 1 over the interval $0 < \tau < 1$ and should be divided by 3 and repeated on the other two edges of the triangle to give the corresponding density over all of $P\mathcal{S}(4, 1)$. A plot of these three densities on the τ scale is given in Fig. 2. Note that f_{Ind} and f_{L2} are difficult to tell apart and have poles at the endpoints $\tau = 0, 1$. In contrast f_{L4} is flat.

It is also of interest to plot these densities on the δ scale, where we rewrite each density as a function of δ and introduce the Jacobian factor $d\tau/d\delta = 2 \sin \delta \cos \delta = \sin 2\delta$ for the change of variables. The resulting densities are in Fig. 3. Now all three densities are quite distinct. Both f_{Ind} and f_{L4} vanish at the endpoints, though f_{L4}

Fig. 2 Densities for three possible “uniform” distributions in τ coordinates, plotted for $0 < \tau < 1$: independence model (*solid*), level 2 uniform (*dashed*) and level 4 uniform (*dotted*)





Fig. 3 Densities for three possible “uniform” distributions in δ coordinates, plotted for $0 < \delta < \pi/2$: independence model (*solid*), level 2 uniform (*dashed*) and level 4 uniform (*dotted*)



converges to 0 more quickly. In the middle of the interval f_{ind} is bimodal whereas f_{L4} is unimodal. In contrast f_{L2} is flat throughout.

6 Constructing Distributions on $P\mathcal{S}(4, 1)$ About a Preferred Shape

Both the spherical and planar representations of $P\mathcal{S}(4, 1)$ are topological circles. Thus one modelling strategy is to ignore the corners and treat them as actual circles. Then fit a standard circular model such as a von Mises distribution, which can allow concentration about any specified projective shape. This strategy was explored by Goodall and Mardia [5].

However, in this paper we look at a different strategy, which is valid for any compact manifold which can be embedded in a Euclidean space. Namely, we construct an exponential family based on first, and possibly second, moments in the Euclidean coordinates. This strategy has been very common and successful in directional data analysis, yielding the Fisher and Bingham distributions and various generalizations.

For the spherical triangle representation of $P\mathcal{S}(4, 1)$, the simplest strategy is to condition the linear-exponential function $\exp(a_1w_1 + a_2w_2 + a_3w_3)$ to lie on the spherical triangle, where w_1, w_2, w_3 are the Euclidean coordinates and a_1, a_2, a_3 are parameters. The resulting distribution will be truncated von Mises on each arc. However, since the cumulative distribution function of the von Mises distribution is a bit awkward to work with, we will not consider this density further here.


For the planar triangle representation of $P\mathcal{S}(4, 1)$, the same strategy of using just first order Euclidean terms in the exponent of the density does not work very well. The resulting densities are not very flexible because the mode of the density must lie either at a vertex or uniformly along an edge.

Hence we consider an exponential family in the plane based on linear and quadratic terms in the Euclidean coordinates. The simplest version of this strategy is to condition an isotropic bivariate normal distribution to lie on the planar triangle. This distribution is explored in the next section.

7 Conditioned Normal Distribution for the Planar Representation of $P\mathcal{S}(4, 1)$

In this section we look at a bivariate normal distribution $N_2(\mu, \sigma^2 I)$, conditioned to lie on the planar equilateral triangle representation of projective shape space.

Fig. 4 The outline of the heavy central triangle is a Level 4 representation of projective shape space for four collinear points. The labels 1–22 demarcate 22 possible regions for the mean parameter μ in the conditioned isotropic bivariate normal distribution



An explicit coordinate representation of this equilateral triangle is given by the choice of vertices

$$v_A = \frac{1}{\sqrt{3}} \begin{bmatrix} -\frac{\sqrt{3}}{2} \\ -\frac{1}{2} \end{bmatrix}, \quad v_B = \frac{1}{\sqrt{3}} \begin{bmatrix} \frac{\sqrt{3}}{2} \\ -\frac{1}{2} \end{bmatrix}, \quad v_C = \frac{1}{\sqrt{3}} \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

so that the edges have length 1, edge AB is horizontal, and the center of the triangle is at the origin.

In order to study the shape of the bivariate normal distribution conditioned to lie on this triangle, it is helpful to divide the parameter space for $\mu \in R^2$ into 22 regions, as shown in Fig. 4. The shape space $P\mathcal{S}(4, 1)$ is denoted by the central triangle with thick edges and vertices marked A,B,C. At each vertex, two lines have been plotted, orthogonal to each of the two edges at the vertex. These lines, plus the original triangle partition R^2 into 22 regions, as marked.

The behavior of the density on a particular edge, e.g. AB, depends on whether μ lies inside or outside the corresponding parallel lines perpendicular to that edge. If μ lies inside the parallel lines, then the quadratic form has a minimum, and hence the density has a local maximum, on the edge. If μ lies outside the parallel lines, then the quadratic form is monotone increasing (and the density is monotone decreasing) along the edge from the nearer parallel line to the further parallel line.

It is convenient to classify the behavior of the conditioned normal density into five types:

- Type I. The density is unimodal with the mode within one edge (regions 4,7,21).
- Type II. The density is unimodal with the mode at a vertex (regions 1,3,8,14, 20,22).

Type III. The density is bimodal with the modes within two edges (regions 5,6,9, 13,16,18).

Type IV. The density is bimodal with the modes at one vertex and within the opposite edge (regions 2,15,19).

Type V. The density is trimodal with a mode within each edge (regions 10,11, 12,17).

Figure 5 gives plots of each of these different types of density, with μ taking the values $-2v_A$, $-2v_A - 4v_B$, $v_A + .2(v_B - v_C)$, $-2v_A - 2v_B$, 0 , respectively. Note that Types I and II are likely to be the most relevant in practice.

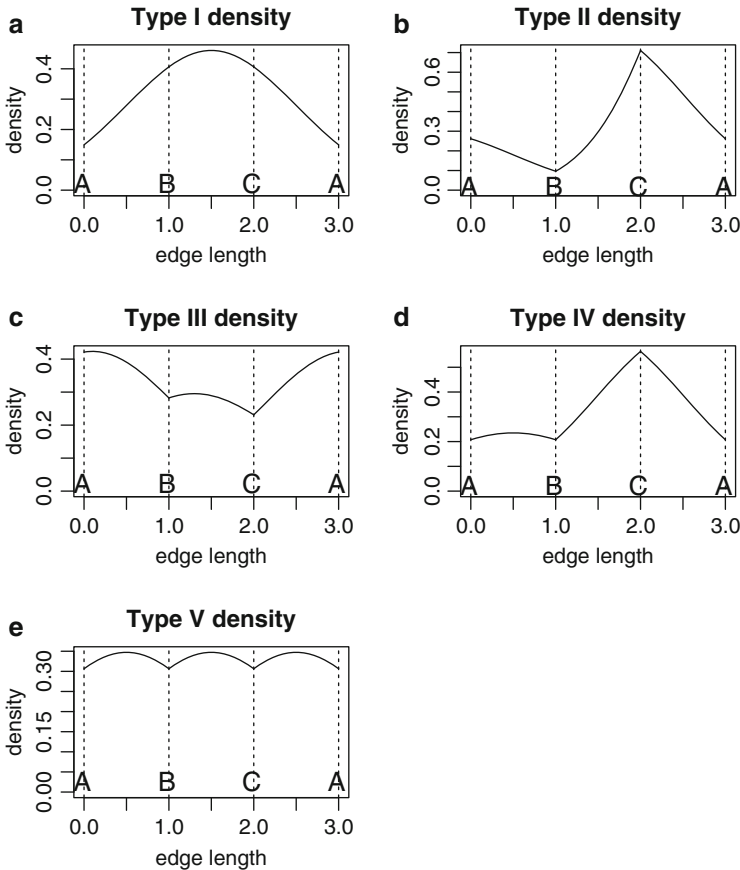


Fig. 5 Plots of the conditioned normal density for the planar triangle representation of $P\mathcal{S}(4, 1)$ along the edges AB, BC, CA. Parts (a)–(e) illustrate densities of Types I–V, respectively

This section has focused on the case of an isotropic covariance matrix $\Sigma = \sigma^2 I$ for the underlying bivariate normal distribution. It is interesting to consider what happens if Σ is unrestricted. In this case it is possible to divide the parameter space for μ into different regions similarly to Fig. 4, except the thin lines are now orthogonal to the edges of the triangle in the Σ metric rather than the Euclidean metric. Provided the level of anisotropy is not too severe, there is still a split into 22 regions. However, if the level of anisotropy is extreme, the thin lines will cross over the edges of the triangle and the partition into regions will be more complicated.

8 Discussion

The Procrustes approach to projective shape offers an elegant mathematical framework to study projective shape. However, the construction of useful densities in this setting is considerably more complicated than in the more traditional setting of similarity shape analysis. In particular, there are at least three plausible candidates for the label of “uniform” distribution on $P\mathcal{S}(4, 1)$ and a variety of approaches to construct more concentrated distributions. More work is needed to fully appreciate the implications of this framework for computer vision and statistical inference.

Acknowledgements This work benefited from a visit to the Statistical and Applied Mathematics Sciences Institute (SAMSI), North Carolina and a workshop at the Mathematical Biosciences Institute (MBI), Ohio, especially through helpful discussions with Ian Dryden, Thomas Hotz, Huiling Le, Stephan Huckemann, Kanti Mardia, Ezra Miller, and Vic Patrangenaru.

References

1. Åström, K., Morin, L.: Random cross ratios. In: Borgefors, G. (ed.) Proceedings of the 9th Scandinavian Conference on Image Analysis, pp. 1053–1061. Swedish Society for Automated Image Analysis, Göteborg (1995)
2. Dryden, I.L., Mardia, K.V.: Statistical Shape Analysis. Wiley, Chichester (1998)
3. Dryden, I.L., Kume, A., Le, H., Wood, A.T.A.: A multi-dimensional scaling approach to shape analysis. *Biometrika* **95**, 779–798 (2008)
4. Faugeras, O., Luong, Q.-T.: The Geometry of Multiple Images. MIT Press, Cambridge (2001)
5. Goodall, C.R., Mardia, K.V.: Projective shape analysis. *J. Comput. Graph. Stat.* **8**, 143–168 (1999)
6. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, Cambridge (2000)
7. Huckemann, S., Hotz, T., Munk, A.: Intrinsic shape analysis: geodesic principal component analysis for Riemannian manifolds modulo isometric Lie group actions (with discussion). *Stat. Sinica* **20**, 1–100 (2010)
8. Kent, J.T., Mardia, K.V.: A geometric approach to projective shape and the cross ratio. *Biometrika* **99**, 833–849 (2012)
9. Maybank, S.J.: Classification based on the cross ratio. In: Mundy, J.L., Zisserman, A., Forsyth, D. (eds.) Applications of Invariance in Computer Vision. Lecture Notes in Computer Science, vol. 825, pp. 453–472. Springer, Berlin (1994)

10. Maybank, S.J.: Probabilistic analysis of the application of the cross ratio to model based vision. *Int. J. Comput. Vis.* **16**, 5–33 (1995)
11. Tyler, D.E.: A distribution-free M -estimator of multivariate scatter. *Ann. Stat.* **15**, 234–251 (1987)
12. Tyler, D.E.: Statistical analysis for the angular central Gaussian distribution on the sphere. *Biometrika* **74**, 579–589 (1987)