

Statistical Shape Analysis of Experiments for Manufacturing Processes

Enrique DEL CASTILLO

Department of Industrial and Manufacturing Engineering
The Pennsylvania State University
University Park, PA 16802
(exd13@psu.edu)

Bianca M. COLOSIMO

Dipartimento di Meccanica
Politecnico di Milano
20133 Milano, Italy

We review existing and develop new statistical techniques for the analysis of experiments where the response is the geometric shape of a manufactured part. The analysis of variance for shapes is discussed. An F test and a permutation test for detecting differences in shape are presented. It is shown how the permutation test provides higher power for 2D circular profiles than the traditional methods used in manufacturing practice, which are based on the circularity form error. The proposed permutation test does not require the error assumptions needed in the F test, which may be restrictive in practice. New visualization tools, including main effect and interaction plots for shapes and deviation from nominal plots are presented to aid in the interpretation of the experimental results. The proposed methods are illustrated with a real manufacturing application in titanium lathe turning.

KEY WORDS: Geometric specifications; Morphometrics; Procrustes methods; Profile data; Registration; Signal response systems.

1. INTRODUCTION

The statistical analysis of geometrical shapes is a relatively new field within the history of statistics. Seminal work on “shape theory” by Kendall (1984) and Bookstein (1986) only appeared in the early 1980s. Shape theory brings statistical analysis to geometry in the sense of Klein (1939) who by the end of the 19th century defined the geometry of an object as those properties which are invariant to certain transformations in a given space. Accordingly, in statistical shape analysis (SSA) the *shape* of an object is defined as all the information of the object that is invariant with respect to similarity transformations in Euclidean space (rotations, translations, and dilations—changes of scale). The goal of SSA is to analyze the shapes of objects in the presence of random error.

Analysis of shapes in manufacturing is critical because geometrical tolerances (specifications) of roundness, flatness, cylindricity, etc., need to be inspected, controlled, or optimized based on a cloud of two-dimensional or three-dimensional measurements taken on the machined surfaces of the part. These tasks are even more complex if the part geometry has a “free form,” that is, there is no standard geometrical construction that can represent the shape, a situation common in advanced manufacturing applications such as in the aerospace sector.

In this article, we review existing and develop new statistical shape analysis techniques for experimental data collected when the response of interest is the shape of a manufactured part. We demonstrate the benefits and potential of statistical shape analysis applied to parts manufactured in a production system. It will be shown that one of the strengths of shape analysis in manufacturing is the ability to work on the space where the objects exist (i.e., we work with coordinate data that retains the geometry of the objects), rather than working on some derived space, for instance, when working with linear combinations of ratios of lengths and angles. Alternatives like MANOVA cannot be applied in general in manufacturing applications given the

typical dimensions of the problem (see Section 3.1). As will be seen, retaining the geometry also has clear advantages with respect to visualization of the results.

A Real Manufacturing Process Example

To illustrate the type of designed experiments we wish to analyze in this article, consider the following machining experiment. A set of 90 titanium alloy (Ti-6Al-4V) specimens was machined by lathe-turning. Two cutting steps were performed to reduce the initial diameter of 20 mm to the final diameter of 16.8 mm (as shown in Figure 1). The original specimens were obtained by vacuum arc remelting followed by forging, rolling, hardening (1 hour at 780°C and then air cooling), and a last phase of centerless grinding. Lathe-turning of the external surface was then performed considering a full factorial 3^2 design, where each of the 9 treatments was replicated 10 times. The two factors under study were A = depth of cut [mm] and B = cutting speed [mm/rev] of the final (finishing) machining step. Values assumed for these two parameters in each of the treatments are shown in Table 1. The machining feed [mm/rev] was selected at specific levels depending on the cutting speed (specifically a feed equal to 0.07, 0.11, and 0.14 mm/rev corresponded to a cutting speed of 80, 70, and 65 m/min, respectively). This type of dependency between the feed and the speed was suggested by the tool supplier in order to keep constant the tool life. For each specimen, the roundness profile was obtained at a fixed distance of 5 mm from the left-hand side of the specimen (shown with a dotted line in Figure 1). The profile was obtained using a coordinate measuring machine (CMM) that measured a set of 64 equally spaced points on each profile. The




Figure 1. The desired geometry of the final specimen obtained by lathe-turning.

goal of the experiment is to determine the effect of depth and cutting speed on the circularity of the parts and to determine the best settings of these factors to achieve the most circular parts. Given that circular shapes are very common in manufacturing, a standard definition of the circularity form error exists in industrial practice (see, e.g., Krulikowsky 1996, who summarizes the ASME standards; ASME 1994). In Section 6, the power to detect differences in circularity for an ANOVA based on this standard definition will be contrasted with the shape analysis methods proposed here. We will return to the analysis of this experiment in Section 5.

An early reference in the analysis of experiments where the responses are shapes is the work by Snee and Andrews (1971). Snee and Andrews's work predated the modern developments in the statistical analysis of shapes, but had the merit of pointing out the importance of experiments where it is necessary to characterize and even optimize the shape of an object. These authors illustrated their ideas with application in agriculture [characterization of the shapes of sweet potatoes as a function of planting date and variety; Snee (1972) studied also the shape of carrots]. These authors used the type of analysis of shape data existing in the pre-SSA era: one analyzes ratios of distances between locations of interest in the object (what today are called the landmarks).

Over the last 30 years, statistical shape analysis techniques have been developed and applied in many areas of the natural sciences where interest is in characterizing differences in shape, for example, biology, paleontology, and geology. A considerable intersection of ideas exist also with image and pattern analysis in computer science. In particular, SSA is known as *geometric morphometrics* in biology, a field in which some authors refer to a "morphometrics revolution" (Adams, Rohlf, and Slice 2004) given the success SSA had over previous techniques used to analyze shapes. For more on the history and foundations

of SSA we refer readers to the book by Dryden and Mardia (1998).

Our interest in shape analysis stems in part from the recent interest in "profile analysis" in the field of statistical process control or SPC (Kang and Albin 2000; Colosimo, Pacella, and Semeraro 2008) (although we do *not* discuss SPC based on shape analysis in this article, this is certainly another potential area of research where SSA ideas can be used). In profile-based SPC, a parametric model is sought that describes the form that the response follows with respect to some variable of interest (in essence, one performs functional data analysis). The parameters of this model are fitted based on process data and then multivariate SPC methods are applied to the estimated parameters.

Another area that shows the relevance of shape analysis in designed experiments in industry is the work on "signal response systems" [see, e.g., the work by Nair, Taam, and Ye (2002) and by Miller and Wu (1996)], where, similar to profile analysis, a given parametric function describes the process behavior. In a second stage, the changes in the fitted parameters are analyzed with respect to the changes in the experimental conditions. The goal is to maintain the process response close to some desired signal (or profile). To date, despite the similarity of this problem with that studied by Snee and Andrews four decades ago, statistical shape analysis techniques have not been used in the analysis of experimental data in industry.

One of the main difficulties when applying profile-based SPC or analyzing signal response experimental data is that in many experiments in industry, the response of interest has a complicated shape (in manufacturing, the geometric shape of the part) for which finding a parametric model is in itself a challenge. By working with the shape directly, SSA techniques avoid the parametric model definition step, allow complicated shapes to be studied, and simplify the presentation of results. In SSA, one works with the whole shape of the object, so the geometry is not "thrown away" (Dryden and Mardia 1998).

The remainder of the article is organized as follows. In Section 2, we review the main ideas of SSA (readers familiar with shape analysis may wish to skip this section). In this section, the notions of shape space, the generalized procrustes algorithm, and tangent space coordinates are discussed. Section 3 presents the main results of this article. A two-way ANOVA for shapes is described. New visualization tools, in particular, main effect shape plots and interaction effect shape plots, are explained. A two-way permutation ANOVA test for shapes useful under less restrictive error assumptions is given in Section 4. The ANOVA tests for differences in shape are illustrated with the titanium alloy lathe turning experiment in Section 5. This includes a discussion of techniques for shape optimization. Finally, Section 6 presents a brief comparative study of the power of different ANOVA tests for detecting effects on circular shapes. The article concludes with a discussion of shape analysis and alternative techniques, including areas of further work.

2. A REVIEW OF SOME STATISTICAL SHAPE ANALYSIS CONCEPTS AND TECHNIQUES

Given that there is a large body of literature on SSA, only the main ideas used in later sections are presented here. For a

Table 1. Factors and levels in the machining experiment

Treatment	A: depth of cut [mm]	B: cutting speed [m/min]
1	0.4	80
2	0.4	70
3	0.4	65
4	0.8	80
5	0.8	70
6	0.8	65
7	1.2	80
8	1.2	70
9	1.2	65

more thorough presentation of SSA we refer readers to Dryden and Mardia (1998), Goodall (1991), Adams, Rohlf, and Slice (2004), Klingenberg and Monteiro (2005), Krim and Yezzi (2006), Davies, Twining, and Taylor (2008) and the recent articles in Srivastava et al. (2010).

In most of SSA, the main goal is statistical inference with shapes, in particular, to test if two objects have equal shapes or not and to analyze shape variability. Some other authors' main interest (e.g., in biology) is to describe how shapes of objects (e.g., species of animals) change with time. In our case, the main goal is to characterize how the shape of a manufactured part is affected as one changes a set of controllable factors in an experiment.

The techniques considered herein are based on shape data obtained by measuring the parts at specific *landmarks*, points of special interest or unique characteristics. Landmarks are points of correspondence that match between objects. A landmark is given by the two-dimensional or three-dimensional cartesian coordinates of a point on the object surface and a given label for the point, usually a sequential number $1, 2, \dots, k$ which corresponds from object to object. Differences in shapes are measured by adding the squared Euclidean distances over all corresponding landmarks of different objects. For an instance in text recognition, when writing the letter "V" three obvious landmark locations are the two end points and the point where the line changes slope. Assignment of landmarks to objects is in itself an important problem; in some areas such as in archeology or biology specific points of the objects are of interest and this assignment is done manually. In manufacturing, considerable amounts of data can be acquired with a coordinate measuring machine (CMM) or through digital images of the objects. If the surface of the part has edges, these could be used to place the landmarks, which should correspond from part to part. If no obvious landmarks exist, there are also automatic procedures that can be used to determine landmarks, for example, based on points of maximum curvature (Mardia 1989). There also exist other SSA approaches not based on landmarks, to which we refer briefly in Section 7. The SSA methods we consider in this article require landmarks already assigned and corresponding between objects in order to solve problems of statistical inference about the shapes of parts.

Landmark-based SSA techniques follow two main steps: first, the objects under consideration are registered or superimposed with respect to each other in order to filter out rotation, translation, and isometric scaling (dilation) effects. This is done because the objects may have different orientations in Euclidean space or have different locations or sizes, and therefore their shapes cannot be initially compared. The main technique for this task is the generalized procrustes algorithm (GPA). Once objects are registered, multivariate statistical methods of inference can be performed on the projections of the shapes on the space tangent to the mean shape, provided the differences between shapes are small. These two steps are explained below. We first give some geometrical notions necessary to understand the algorithms.

2.1 Preshape and Shape Space

Let \mathbf{X} be a $k \times m$ matrix containing the k landmarks (coordinate pairs or triples) of an object in m (two or three) dimensions. \mathbf{X} is sometimes called a *configuration matrix* (since it is

an element of the configuration space, the space of all possible arrangements of k landmarks in m dimensions), which we could also refer to as a "profile matrix," following manufacturing practice for the case of 2D closed contours (ASME 1994). With this notation, the shape of a configuration \mathbf{X} is obtained, first, by removing location and scale effects by computing the so-called *preshape* \mathbf{Z} :

$$\mathbf{Z} = \frac{\mathbf{H}\mathbf{X}}{\|\mathbf{H}\mathbf{X}\|}, \quad (1)$$

where \mathbf{H} is a $(k-1) \times k$ Helmert submatrix (Dryden and Mardia 1998) and $\|\cdot\|$ denotes the Frobenius norm of a matrix. If we define $h_j = -[j(j+1)]^{-1/2}$, then \mathbf{H} is a matrix whose j th row is

$$(\underbrace{h_j, h_j, \dots, h_j}_{j \text{ times}}, \underbrace{0, \dots, 0}_{k-j-1 \text{ times}}, -jh_j) \quad \text{for } j = 1, \dots, k-1.$$

Note that $\mathbf{H}\mathbf{H}' = \mathbf{I}_{k-1}$ and that the rows of \mathbf{H} are contrasts. Alternatively, one could start with the *centered preshapes*, defined by $\mathbf{Z}_c = \mathbf{H}'\mathbf{Z}$ (these are $k \times m$ matrices).

Transformation (1) removes location effects via the numerator, and rescales the configurations to unit length via the denominator. Since we have not removed rotations from \mathbf{Z} it is not yet the shape of \mathbf{X} , hence the name *preshape*. The centered preshapes are equivalent to centering each coordinate of each configuration by its centroid and dividing each by its norm.

The *shape* of configuration \mathbf{X} , denoted $[\mathbf{X}]$, is defined as the geometrical information that is invariant to similarity transformations. Once location and scale effects are filtered as above, the shape is then defined as

$$[\mathbf{X}] = \{\mathbf{Z}\mathbf{\Gamma} : \mathbf{\Gamma} \in \text{SO}(m)\}, \quad (2)$$

where \mathbf{Z} is the *preshape* of \mathbf{X} , $\mathbf{\Gamma}$ is a rotation matrix [i.e., a matrix such that $\mathbf{\Gamma}'\mathbf{\Gamma} = \mathbf{\Gamma}\mathbf{\Gamma}' = \mathbf{I}_m$ with $\det(\mathbf{\Gamma}) = +1$] and $\text{SO}(m)$ is the space of all $m \times m$ rotation matrices that exclude reflections, the special (or nonreflective) orthogonal group. Multiplication by a suitable matrix $\mathbf{\Gamma}$ reorients (rotates) the object. Note that a shape is therefore defined as a set.

The following geometrical interpretation of these transformations is due to Kendall (1984 and 1989). Given that preshapes are scaled and centered objects, they can be represented by vectors from the center to the surface of a unit sphere of dimension $(k-1)m$, because the numerator in (1) removes m degrees of freedom for location parameters and the denominator removes one additional degree of freedom for the change of scale. The preshapes, having unit length, form a space (denoted S_m^k), which has $(k-1)m-1$ dimensions by virtue of being on the surface. As one rotates a preshape \mathbf{Z} via (2), the vectors $\mathbf{Z}\mathbf{\Gamma}$ describe an *orbit* on S_m^k . All the vectors on an orbit correspond to the same shape, since by definition the shape of an object is invariant to rotations. Thus, the orbits (also called *fibers*) of the preshape space are mapped one to one into single points in the *shape space* (denoted Σ_m^k), the space of all possible shapes of k landmarks in m dimensions. This space in general will be a non-Euclidean M -dimensional manifold. Two objects have the same shape if and only if their preshapes lie on the same fiber. The shape space has dimension $M = (k-1)m-1-m(m-1)/2$ since in addition to losing location and dilation degrees of freedom we also lose $m(m-1)/2$ degrees of freedom in the specification of the (symmetric) $m \times m$ rotation matrix $\mathbf{\Gamma}$.

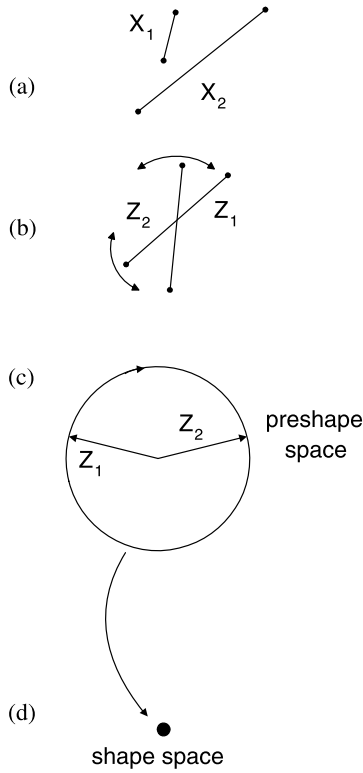


Figure 2. One of the simplest illustrations of preshape and shape space. (a) Two lines in the original two-dimensional space; (b) preshapes on two-dimensional Euclidean space, after centering and scaling; (c) the corresponding preshape space is the (one-dimensional) circumference of a unit circle. The two preshapes lie on the single fiber or orbit generated as the preshapes are rotated, hence there is a single shape; (d) the shape space for the two lines (Σ_2^2) is zero dimensional (a single point) and corresponds to the only shape that exists in this example.

Example (Preshape space and shape space). In order to explain these ideas, consider one of the simplest possible cases, where we have two lines in \mathbb{R}^2 (see Figure 2). Thus, we have that $m = 2$ and $k = 2$, where the obvious landmarks are the end-points of the lines. After centering and scaling the two lines using (1), one obtains the preshapes with matrices \mathbf{Z}_1 and \mathbf{Z}_2 . Since the original objects evidently have the same shape (that of a line in Euclidean space) these two preshapes lie on the same fiber or orbit, generated as the preshapes are rotated using (2). The preshape space S_2^2 is of dimension $(k-1)m-1 = 1$, namely, the circumference of a unit circle. As the preshapes rotate (they can rotate clockwise or counterclockwise) they will eventually coincide, which corresponds to the centered and scaled lines coinciding. Finally, since there is a single shape, the shape space Σ_2^2 is the simplest possible, namely, a single point [dimension is $M = (k-1)m-1-m(m-1)/2 = 0$, i.e., a 0-manifold].

In general, the shape space Σ_m^k will be a nonlinear space, the Riemannian M -manifold formed by the landmarks modulo similarity transformations, of reduced dimension than the always spherical preshape space. That is, the shape space is defined as a quotient space, that is, $\Sigma_m^k = \mathbb{R}^{km}/\mathcal{G} = S_m^k/\text{SO}(m)$, where \mathcal{G} is the group of similarity transformations. While the step of going from configuration space (the km -manifold of all possible

arrangements of the landmarks) to preshape space is easy to understand, going from preshape space to shape space is a nontrivial step. For instance, for planar shapes Kendall (1984) showed that $\Sigma_2^k = \mathbb{C}P^{k-2}(4)$, the complex projective space of sectional curvature 4 [thus in the previous example, $\Sigma_2^2 = \mathbb{C}P^0(4)$, a one-point space]. See Kendall et al. (1999) for a detailed discussion of the topology of shape spaces. Fortunately for applications in manufacturing, the shapes will typically be very close in shape space, and therefore the nonlinearity of the space can be neglected.

2.2 Generalized Procrustes Algorithm

Two preshapes \mathbf{Z}_1 and \mathbf{Z}_2 lying on different fibers correspond to two objects with different shapes. A measure of the similarity between two shapes is the shortest distance between the fibers, the *procrustes distance* $\rho(\mathbf{X}_1, \mathbf{X}_2)$. This corresponds to the shortest distance along the surface of the preshape space and is therefore a distance along a geodesic. Alternatively, two measures of distance over a linear space are the “partial procrustes distance,” given by

$$d_p(\mathbf{X}_1, \mathbf{X}_2) = \min_{\Gamma \in \text{SO}(m)} \|\mathbf{Z}_2 - \mathbf{Z}_1 \Gamma\| \quad (3)$$

and the “full procrustes distance,” where the minimization is also done over a scale parameter:

$$d_F(\mathbf{X}_1, \mathbf{X}_2) = \min_{\Gamma \in \text{SO}(m), \beta \in \mathbb{R}} \|\mathbf{Z}_2 - \beta \mathbf{Z}_1 \Gamma\|. \quad (4)$$

Geometrically, $d_p(\mathbf{X}_1, \mathbf{X}_2)$ is the secant between \mathbf{Z}_1 and \mathbf{Z}_2 in preshape space, and $d_F(\mathbf{X}_1, \mathbf{X}_2)$ is the distance along the tangent at either one of the preshapes (see Figure 3). As can be seen, for objects with similar shapes, $\rho \approx d_F \approx d_p$. Note that all these metrics are extrinsic to the shape space.

For a collection of n registered configurations or profiles, the *generalized procrustes algorithm* registers or superimposes all




Figure 3. Distances between two shapes in preshape space. ρ is the procrustes distance (along a geodesic), d_F is the full procrustes distance (along a tangent), and d_p is the partial procrustes distance (along the secant). The preshapes have $\|\mathbf{Z}_i\| = 1$, and the preshape space will be, in general, a unit hypersphere of $(k-1)m-1$ dimensions.

the n objects by finding scaling factors $\beta_i \in \mathbb{R}$, rotation matrices $\Gamma_i \in \text{SO}(m)$ and m dimensional translation vectors γ_i , $i = 1, \dots, n$, such that they minimize the sum of squared full procrustes Distances between all objects:

$$G(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n) = \min_{\beta_i, \Gamma_i, \gamma_i} \frac{1}{n} \sum_{i=1}^n \sum_{j=i+1}^n \|\beta_i \mathbf{X}_i \Gamma_i + 1_k \gamma_i' - (\beta_j \mathbf{X}_j \Gamma_j + 1_k \gamma_j')\|^2, \quad (5)$$

where 1_k is a vector of k ones. Constraints must be added to avoid the trivial solution where all parameters are zero (see below). The resulting registered configurations are called the *full procrustes fits*, defined as

$$\mathbf{X}_i^p = \hat{\beta}_i \mathbf{X}_i \hat{\Gamma}_i + 1_k \hat{\gamma}_i', \quad i = 1, \dots, n. \quad (6)$$

The mean shape of the n objects is simply the average of the n registered configurations, namely, $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i^p$.

The minimization (5) needs to be subjected to a constraint that limits the scaling done, otherwise the optimal value of G will be zero. One such restriction is to use a constraint on the size of the mean shape: $S(\hat{\mu}) = 1$ where the size of any configuration \mathbf{X} is defined as $S(\mathbf{X}) = \sqrt{\sum_{i=1}^k \sum_{j=1}^m (X_{ij} - \bar{X}_j)^2} = \|\mathbf{C}\mathbf{X}\|$, where $\mathbf{C} = \mathbf{I}_k - k^{-1} 1_k 1_k'$, $\bar{X}_j = \frac{1}{n} \sum_{i=1}^k X_{ij}$, and X_{ij} is the j th coordinate of the i th point in the configuration. Another common constraint, used in what follows, is to make the average of the squared sizes of the registered configurations \mathbf{X}_i^p given by (6) equal to the average of the squared sizes of the original objects:

$$\frac{1}{n} \sum_{i=1}^n S^2(\mathbf{X}_i^p) = \frac{1}{n} \sum_{i=1}^n S^2(\mathbf{X}_i). \quad (7)$$

The generalized procrustes algorithm, as developed by Gower (1975) and Ten Berge (1977) proceeds as follows to solve (5) subject to (7):

1. Center (but do not scale) the configurations $\mathbf{X}_1, \dots, \mathbf{X}_n$ by initially defining

$$\mathbf{X}_i^p = \mathbf{H}\mathbf{X}_i, \quad i = 1, \dots, n$$

[alternatively, we can define $\mathbf{H}'\mathbf{H}\mathbf{X}_i = \mathbf{C}\mathbf{X}_i = \mathbf{X}_i^p$ and the resulting matrices will be $k \times m$; note that \mathbf{X}_i^p as defined above is instead $(k-1) \times m$].

2. Let $\bar{\mathbf{X}}_{(i)} = \frac{1}{n-1} \sum_{j \neq i} \mathbf{X}_j^p$, $i = 1, \dots, n$. These are the “jackknifed” average shapes excluding object i .
3. Do a procrustes fit (rotation only) of the current \mathbf{X}_i^p 's on to $\bar{\mathbf{X}}_{(i)}$. This yields rotation matrices $\hat{\Gamma}_i$ from which we let

$$\mathbf{X}_i^p \leftarrow \hat{\Gamma}_i \mathbf{X}_i^p, \quad i = 1, \dots, n.$$

We repeat steps 2 and 3 for all i .

4. Compute the $n \times n$ correlation matrix $\Phi = \text{corr}(\mathbf{X}_v)$, where

$$\mathbf{X}_v = [\text{vec}(\mathbf{X}_1^p) \text{vec}(\mathbf{X}_2^p) \dots \text{vec}(\mathbf{X}_n^p)].$$

Note we stack all the m dimensions together.

5. Let $\phi = (\phi_1, \dots, \phi_n)'$ be the eigenvector of Φ corresponding to its largest eigenvalue. Then set

$$\hat{\beta}_i = \sqrt{\frac{\sum_{j=1}^n \|\mathbf{X}_j^p\|^2}{\|\mathbf{X}_i^p\|^2}} \phi_i, \quad i = 1, \dots, n,$$

and let $\mathbf{X}_i^p \leftarrow \hat{\beta}_i \mathbf{X}_i^p$. The algorithm repeats steps 2 to 5 until convergence.

The algorithm is guaranteed to converge, and converges usually in just a few iterations (Ten Berge 1977). The exact solution to the procrustes registration problem between two objects \mathbf{X}_1 and \mathbf{X}_2 required in step 3, implies finding $\Gamma \in \text{SO}(m)$ that minimizes $d_p(\mathbf{X}_1, \mathbf{X}_2)$ [see Equation (3)] for $\mathbf{X}_1 = \mathbf{X}_i^p$ and $\mathbf{X}_2 = \bar{\mathbf{X}}_{(i)}$, $i = 1, \dots, n$. The exact solution to this problem is well known in both statistics (Jackson 2003) and computer vision (Horn, Hilden, and Negahdaripour 1988) and is given by $\hat{\Gamma} = \mathbf{U}\mathbf{V}'$ where \mathbf{U} and \mathbf{V} are obtained from the singular value decomposition $\mathbf{Z}_2' \mathbf{Z}_1 = \mathbf{V}\mathbf{A}\mathbf{U}$. An important implementation detail of singular value decomposition for shape analysis is that to assure we have $\det(\hat{\Gamma}) = +1$ and hence a rotation matrix (as opposed to -1 and a reflection matrix), we can make instead $\hat{\Gamma} = \mathbf{U}\mathbf{R}\mathbf{V}'$ where \mathbf{R} is the identity matrix except for its last diagonal which equals $\det(\mathbf{U}\mathbf{V}')$.

The GPA algorithm as described assumes the statistical model

$$\mathbf{X}_i = \beta_i(\mu + \mathbf{E}_i)\Gamma_i + 1_k \gamma_i', \quad i = 1, \dots, n, \quad (8)$$

where μ is the mean configuration and the $k \times m$ matrix of errors \mathbf{E}_i is such that $\text{vec}(\mathbf{E}_i) \sim (\mathbf{0}, \sigma^2 \mathbf{I})$, where $\mathbf{0}$ is a vector of km zeroes and \mathbf{I} is the $km \times km$ identity. The model then assumes *isotropic variance*, that is, the variance is the same at each landmark and in each of the m coordinates. Modification of GPA for the case of a general covariance matrix of the errors Σ requires a straightforward modification of the definition of the d_F distances minimized in (5) that accounts for Σ . However, given that in general Σ is unknown and needs to be estimated there is no known registration algorithm which guarantees convergence in the nonisotropic case. Common practice is to initially set $\Sigma = \mathbf{I}$, run GPA, then estimate Σ with

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \text{vec}(\mathbf{X}_i^p - \hat{\mu}) \text{vec}(\mathbf{X}_i^p - \hat{\mu})',$$

run again GPA with the squared full procrustes distances in (5) replaced by the Mahalanobis squared distance $\text{vec}(\mathbf{X}_i^p)' \hat{\Sigma}^{-1} \times \text{vec}(\mathbf{X}_j^p)$ and iterate this process (but convergence is not guaranteed).

Model (8) implies that each object results from the rotation, scaling, and translation of the mean shape in the presence of random noise, that is, similarity transformations of the mean shape observed with noise generate the observed profiles of the objects.

2.3 Tangent Space Coordinates

Once n configurations or profiles have been registered using GPA, the mainstream of the SSA literature (see, e.g., Dryden and Mardia 1998; Goodall 1991; Adams, Rohlf, and Slice

2004) recommends that further statistical analysis of shape variability and any desired inferences be made based on the resulting registered shapes \mathbf{X}_i^p using the full procrustes distances from the mean shape (or *pole*), called the *tangent space coordinates*. This is appropriate if the shapes are close to each other, and hence instead of considering distances in the nonlinear shape space we can approximate these with the linear tangent space distances (if shapes vary considerably, see Huckemann, Hotz, and Munk 2010a, for a recent MANOVA on the nonlinear shape space). A principal component analysis (PCA) is also recommended on the tangent space coordinates to better understand the directions in which the shapes are varying the most [the corresponding PCA on the nonlinear shape space has been studied by Huckemann, Hotz, and Munk (2010b), for cases when between shape variability is large].

Performing a PCA on the tangent coordinates is of value when one is interested in analyzing how the *variability* of the shapes behaves around the mean shape. For analyzing the effect of factors (varied during an experiment) on the *mean* shape (as we require in the present article) one needs to perform an analysis of variance (ANOVA), which we now discuss.

3. ANALYSIS OF VARIANCE OF SHAPES

Goodall (1991) (see also Dryden and Mardia 1998) shows that an approximate balanced one-way ANOVA can be performed to test for the equality of mean shape among a groups of n objects if the shapes are close to each other. The more general unbalanced one-way ANOVA decomposition with a treatments is given in terms of the full procrustes distances, that is,

$$\begin{aligned} SS_{total} &= \sum_{i=1}^a \sum_{j=1}^{n_i} d_F^2(\mathbf{X}_{ij}^p, \bar{\mathbf{X}}_{..}) \\ &\approx \sum_{i=1}^a n_i d_F^2(\bar{\mathbf{X}}_{i.}, \bar{\mathbf{X}}_{..}) + \sum_{i=1}^a \sum_{j=1}^{n_i} d_F^2(\mathbf{X}_{ij}^p, \bar{\mathbf{X}}_{i.}) \\ &= SS_{treatments} + SS_{error}, \end{aligned}$$

where as usual in the balanced case $n_i = n$ for all i and $SS_{treatments} = n \sum_{i=1}^a d_F^2(\bar{\mathbf{X}}_{i.}, \bar{\mathbf{X}}_{..})$. These are only approximate expressions, becoming closer to an equality as the error variance decreases and as the shapes get closer to the mean shape. A diagrammatic depiction of the terms in this expression is shown (on preshape space) on Figure 4 for the case of two treatments and two objects.

From Figure 4, and the relations $\sin^2 \rho = d_F^2(\mathbf{X}_1, \mathbf{X}_2)$ (note that the procrustes distance ρ can be considered an angle, see Figure 3) and $\rho_3 = 2\rho_2 - \rho_1$, we have that

$$\begin{aligned} SS_{total} &= 2(\sin^2(2\rho_2 - \rho_1) + \sin^2 \rho_1) \approx SS_{treatments} + SS_{error} \\ &= 4\sin^2 \rho_2 + 4\sin^2(\rho_2 - \rho_1). \end{aligned}$$

In general, we have that $SS_{total} \leq SS_{treatment} + SS_{error}$. This gets closer to an equality the smaller ρ_2 gets, which occurs when the treatment mean shapes get closer to the overall mean shape. In addition, the departure from equality tends to 0 as $\rho_2 - \rho_1$ tends to 0, which indicates the case when the differences around the treatment mean (given by the error term variance, σ^2), reduce to zero.




Figure 4. One-way ANOVA illustration, $a = 2, n_1 = n_2 = n = 2$. The sum of the squares of the depicted distances enters into the ANOVA expression. Only when all shapes are close to the overall mean (much closer than depicted) the curvature of the shape space can be neglected and the ANOVA identity will be approximately true. This is usually the case in the type of manufacturing applications that concern us here.

Applying results from Langron and Collins (1985, theorem 6.2) we have that for the isotropic, normal model (8), where $M = (k - 1)m - 1 - m(m - 1)/2$ is the dimension of the shape space,

$$\sum_{i=1}^a n_i d_F^2(\bar{\mathbf{X}}_{i.}, \bar{\mathbf{X}}_{..}) \sim \sigma^2 \chi_{(a-1)M}^2 \quad (9)$$

and

$$\sum_{i=1}^a \sum_{j=1}^{n_i} d_F^2(\mathbf{X}_{ij}^p, \bar{\mathbf{X}}_{i.}) \sim \sigma^2 \chi_{(\sum_{i=1}^a n_i - a)M}^2, \quad (10)$$

with the first result above holding only under $H_0: [\mu_1] = [\mu_1] = \dots = [\mu_a] = [\mu]$. From these results, Goodall (1991) suggests to test H_0 using

$$F_0 = \frac{SS_{treatments}}{(a-1)M} \bigg/ \frac{SS_{error}}{(N-a)M} \quad (11)$$

which for small σ follows a $F_{(a-1)M, (N-a)M}$ distribution under H_0 . The isotropic model assumptions (normal errors with constant variance in each dimension and no correlation between dimensions or between landmarks) are quite restrictive. For example, quite often in practice the two coordinates of a landmark are correlated. In consequence, we suggest instead to test H_0 using a permutation test (see Section 4) which holds true under more general conditions.

3.1 Two-Way ANOVA of Shapes

Some authors have considered more complex analysis of variance models for the study of symmetry in biology (Klingenberg and McIntyre 1998; Klingenberg and Monteiro 2005). In manufacturing applications, it is of utmost importance to be able to characterize the effect of controllable factors on the shape of parts. For this purpose, Goodall's one-way ANOVA model can be extended to a factorial-like structure in a straightforward manner. For a two-way balanced ANOVA for shapes, the mean of the object configurations is identical to a standard balanced two-way model:

$$\begin{aligned} E[\mathbf{X}_{ijl}] &= \mu + \tau_i + \beta_j + (\tau\beta)_{ij}, \\ i &= 1, \dots, a; j = 1, \dots, b; l = 1, \dots, n, \end{aligned} \quad (12)$$

where the different effect matrices have the same dimensions as \mathbf{X}_{ijl} . We consider only balanced ANOVAs in the rest of the article; a discussion on the difficulties found with unbalanced designs in two-way ANOVA is given in Section 7. The error terms are assumed to follow the isotropic error variance model, that is, $\epsilon_{ijl} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ (iid). We point out that a MANOVA cannot be performed on the vectorized \mathbf{X}_{ijl} matrices given that we usually have in practice that $km \geq ab(n-1)$ (see Press 2005, p. 267), owing to the large number of measurements per part available in manufacturing.

Since the parts may not be initially registered, the first step of the ANOVA on shapes is to register all $abn = N$ shapes with the GPA algorithm. The procrustes fits will then be given by $\mathbf{X}_{ijl}^p = \hat{\beta}_{ijl} \mathbf{X}_{ijl} \hat{\Gamma}_{ijl} + 1_k \gamma'_{ijl}$. The estimated overall mean shape is $\hat{\mu} = \bar{\mathbf{X}}... = 1/N \sum_i \sum_j \sum_l \mathbf{X}_{ijl}^p$ where we use the standard “dot” notation in two-way ANOVA. We work with the registered profiles \mathbf{X}_{ijl}^p from then on.

Provided the shapes of the objects are not too far from the mean shape (a condition that should hold for each of the $ab + a + b + 1$ means in a two-way ANOVA), we have the approximate ANOVA partition

$$SS_{total} \approx SS_A + SS_B + SS_{AB} + SS_{error},$$

where $SS_{total} = \sum_{i=1}^a \sum_{j=1}^b \sum_{l=1}^n d_F^2(\mathbf{X}_{ijl}^p, \bar{\mathbf{X}}...)$,

$$SS_A = bn \sum_{i=1}^a d_F^2(\bar{\mathbf{X}}_{i..}, \bar{\mathbf{X}}...),$$

$$SS_B = an \sum_{j=1}^b d_F^2(\bar{\mathbf{X}}_{.j.}, \bar{\mathbf{X}}...),$$

$$SS_{AB} = n \sum_{i=1}^a \sum_{j=1}^b d_F^2(\bar{\mathbf{X}}_{ij.} - (\bar{\mathbf{X}}_{i..} - \bar{\mathbf{X}}...) - (\bar{\mathbf{X}}_{.j.} - \bar{\mathbf{X}}...), \bar{\mathbf{X}}...)$$

(note how in SS_{AB} the row and column effects are discounted in the first argument of the distance formula), and where

$$SS_{error} = \sum_{i=1}^a \sum_{j=1}^b \sum_{l=1}^n d_F^2(\mathbf{X}_{ijl}^p, \bar{\mathbf{X}}_{ij.}).$$

Given that the ANOVA is valid only if shapes are close to each other (otherwise the nonlinearity of the shape space must be considered), an important practical question is: when can we claim the sizes of objects are “close enough” to the mean so that the ANOVA approximation holds? Dryden and Mardia (1998, p. 287) give the following rule of thumb: objects are close to the mean shape if $\max d_F < 0.2$. (Note this does not consider the number of landmarks, k , so it should be seen only as a rough recommendation. In the type of manufacturing applications we consider, $\max d_F$ is much smaller than 0.2 despite k being usually in the hundreds. See Section 5 for an example.)

An extension of the one-way ANOVA in Goodall (1991) for testing $H_0^{(1)}: \tau_i = \mathbf{0}$, $H_0^{(2)}: \beta_j = \mathbf{0}$, and $H_0^{(3)}: (\tau\beta)_{ij} = \mathbf{0}$ is based on the statistics

$$F_0^{(1)} = MS_A / MS_{error},$$

$$F_0^{(2)} = MS_B / MS_{error},$$

$$F_0^{(3)} = MS_{AB} / MS_{error}.$$

Similarly to the one-way ANOVA case, the distributions of these three statistics (under their null hypothesis for small σ) are

$$F_{(a-1)M, ab(n-1)M}, \quad F_{(b-1)M, ab(n-1)M}, \quad \text{and} \\ F_{(a-1)(b-1)M, ab(n-1)M},$$

respectively. But just as for the one-way ANOVA case, given the restrictive assumptions in the isotropic model, the permutation tests in Section 4 are recommended instead.

The estimated effects are preshapes, that is, $(k-1) \times m$ normalized configuration matrices (or $k \times m$ matrices, if centered). The algebraic form of the least squares estimates is identical to the ones used in standard ANOVA, and requires additional constraints to be uniquely obtained. Under the usual constraints [i.e., $\sum_{i=1}^a \tau_i = \mathbf{0}$, $\sum_{j=1}^b \beta_j = \mathbf{0}$, $\sum_{i=1}^a (\tau\beta)_{ij} = \mathbf{0}$, and $\sum_{j=1}^b (\tau\beta)_{ij} = \mathbf{0}$; Searle 1971] we have that

$$\hat{\tau}_i = \bar{\mathbf{X}}_{i..} - \bar{\mathbf{X}}..., \quad (13)$$

$$\hat{\beta}_j = \bar{\mathbf{X}}_{.j.} - \bar{\mathbf{X}}..., \quad (14)$$

$$(\hat{\tau\beta})_{ij} = \bar{\mathbf{X}}_{ij.} - \bar{\mathbf{X}}_{i..} - \bar{\mathbf{X}}_{.j.} + \bar{\mathbf{X}}..., \quad (15)$$

where all means are computed from the registered profiles.

3.2 Visualization of ANOVA Effects: Shape Main Effect and Shape Interaction Plots

It is difficult in general to interpret the effect estimates (13)–(15) directly. However, one of the main advantages of the proposed approach is that we preserve the geometry of the objects, and hence visualization of the multivariate results is in principle simple and useful. Visualization to display the directions and magnitude of the *variability* of the landmarks based on principal component analysis is frequently used in the SSA literature (see, e.g., Dryden and Mardia 1998). Here we propose analogous plots to display the directions and magnitudes of the *effects* of the factors on the mean shape. While it is tempting to plot the average shape at each factor combination ($\bar{\mathbf{X}}_{ij.}$), with modern manufacturing data this will almost never help, since the differences will be of such a small order to be undetectable by the human eye (“zooming in” of several orders of magnitude will be required to see differences, but this will lose the overall geometric feature one wishes to observe). Instead, we suggest to plot the effects (13)–(15) at each landmark in a scaled form to allow visualization. We define a *shape main effect* in the most obvious way, namely, as the change in mean shape when a factor is changed. To visualize these effects, we need to refer these changes to the mean shape. Thus, if we simply plot $\hat{\tau}_i$ or $\hat{\beta}_j$, we will not see the mean shape of the objects, as the effects (13)–(15) are differences between average shapes.

In a main effect shape plot we first graph points giving the overall mean shape $\hat{\mu} = \bar{\mathbf{X}}...$. At each landmark of the mean shape, we then add arrows proportional to the corresponding coordinates in the main effect matrix. Thus, for factor A, we add arrows proportional to the corresponding coordinates of $\hat{\tau}_i$. We do this for $i = 1, \dots, a$, showing these plots one on top of each other with different colors for each i , and do likewise to visualize the effect of factor B (see, e.g., Figure 5). If the magnitude of the effects is small relative to the mean shape (not uncommon in




Figure 5. Main effect on shapes for factors A (a) and B (b). Each factor was varied at two levels only. The continuous line corresponds to the sample mean shape $\bar{\mathbf{X}}_{...} = \hat{\boldsymbol{\mu}}$. The effect of A is to induce an oval deformation (a bilobe) that alternates direction from low to high setting of A. As factor B increases, it increases the depth of the notch.

manufacturing), it is necessary to exaggerate the effects by multiplying by a suitable large number. We have found vector plots (made with the `quiver` command in Matlab) especially useful for this purpose. The resulting shape main effect plots have therefore an interpretation equivalent to that of the usual main effect plots given by most DOE software. These plots preserve the geometrical information of the effects, displaying at each landmark the direction in which the average shape changes as the factor is varied.

Visualizing shape interaction effects requires more care if an interaction plot on the shapes equivalent to that of standard DOE interaction plots is desired. Recall how in standard DOE a two-factor interaction plot is constructed (see, e.g., Montgomery 2006). These graphs are prepared not by plotting the interaction effects $(\tau\beta)_{ij}$ directly, but by plotting the *conditional* effects of B given A (or vice versa) for the different combinations of levels for A and B. These effects are computed using “cell” means. Thus, in standard DOE, we conclude—informally—that there is an AB interaction when the effect of A as it changes from its low setting to its high setting *depends* on the setting of factor B (and vice versa). In such case the lines in the interaction plot—which join the averages corresponding to the same level of a factor—will not be parallel.

To visualize interaction effects on the shape of the objects, we need to refer the effects to the mean shape. Thus, similarly as for shape main effects, we first plot the overall mean shape $\hat{\boldsymbol{\mu}} = \bar{\mathbf{X}}_{...}$. We do this b times (one per level of factor B) forming an array of b graphs to be displayed together. At each landmark of the mean shape of the j th graph ($j = 1, \dots, b$) we plot arrows proportional to the corresponding coordinates in the matrix

$$\bar{\mathbf{X}}_{ij.} - \bar{\mathbf{X}}_{...}, \quad i = 1, \dots, a.$$

Visualization is aided if the a arrows at the mean shape landmarks are given in different colors. This gives a visual representation of the conditional effect of B given A. One can alternatively plot an array of a graphs next to each other each with

b arrows in different colors to show the effect of A given B (giving an equivalent “BA” shape interaction plot).

The interpretation of these plots is done by comparing the AB (or BA) shape interaction effect plot with either the A (respectively, the B shape) main effect plot: if the effect of A when changed across its levels is not affected by the setting used on B (so the effect is similar as depicted in the factor A main effect shape plot), we can conclude, informally, that there is no AB interaction effect on the shapes of the objects. For two examples of these plots, see Figure 6.

Example (Shape main effect and shape interaction effect plots). To illustrate the visualization tools described above, we simulated $n = 10$ profiles at each of $2(= a = b)$ levels of two hypothetical factors affecting the mean shape depicted in Figure 7. This type of “notched” circular profile is common in parts that are used for coupling mechanical components. This is a “free-form” of relatively complex geometry not easily described with standard methods. The first factor (A) was simulated to induce a bilobed shape (oval) around the mean; the second factor (B) was simulated to affect the depth of the notch section. Initially, no AB interaction was simulated. Independent $N(0, 0.4^2)$ errors were added to each of the cartesian coordinates of each landmark in the mean shapes. All numerical results in this article were conducted using Matlab (release 2007) and its built-in random number generator (`randn`).

Figure 5 shows the main effect shape plots for factors A and B. Figure 6 shows two AB interaction plots, for the case where no AB interaction was simulated and also for the case where such interaction was simulated.

3.3 ANOVA Test on the Size of the Objects

The ANOVA tests shown in this article deal with shapes, and therefore no consideration is made of the sizes of the objects. While it is possible to modify the generalized procrustes algorithm to consider only rotation and location effects, a simpler

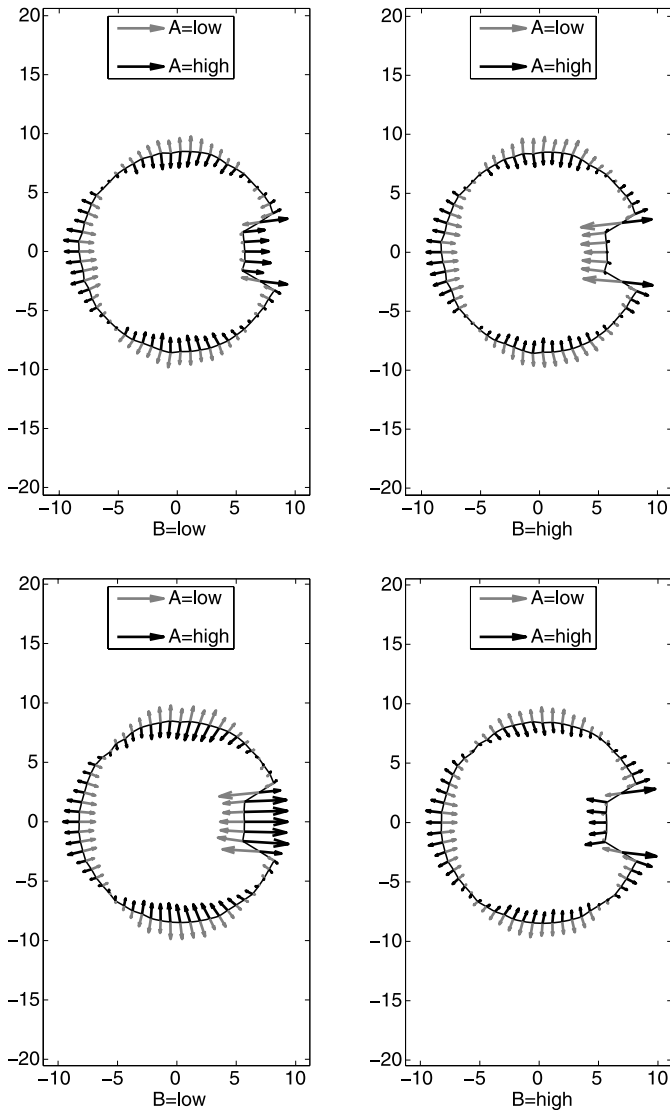


Figure 6. Examples of AB shape interaction plots. On top, the case where no AB interaction exists: changing factor A results in a bilobed shape regardless of the setting used for factor B [compare with Figure 5(a)] while factor B seems the only factor affecting the depth of the notch section. On the bottom: the case where an AB interaction exists. Here, the depth of the notch section is affected by both A and B factors. Factor B no longer determines the depth of the notch section alone.

approach (Dryden and Mardia 1998) that preserves the ANOVA analysis on shapes just presented is to consider in addition a two-way ANOVA test specifically designed to detect changes in size due to the factors. The sizes of the $N = abn$ objects are computed from $S(\mathbf{X}_{ijl})$ (see Section 2.2) and a standard ANOVA is conducted on these sizes. This test is also illustrated in Section 5.

4. A PERMUTATION TEST FOR THE TWO-WAY ANOVA ON SHAPES

As discussed in Section 3, the F_0 statistics used in the ANOVA for shapes hold under restrictive assumptions. Instead, we recommend using the following permutation test.




Figure 7. True mean overall shape (μ), used in the simulated manufacturing example.

A permutation test for the difference between the mean shape of two groups of objects was suggested by Dryden and Mardia (1998). In this spirit, we can also perform a permutation test for the two-way ANOVA case as an alternative to the F tests.

Following recommended practice in random permutation tests (Edgington and Onghena 2007; see also Klingenberg and McIntyre 1998), we rearrange the $N = abn$ objects forming different random subgroups to perform the different tests (for the main effects and the two-factor interaction). While testing for a main effect we account for the assignment of levels for the other factor not being tested. Thus, we form several random arrangements of a subgroups (each with $N/a = bn$ objects) to test for the hypothesis of no difference in mean shapes across rows (factor A). To form these arrangements, we only permute shapes between the levels of factor A but not between the levels of factor B, that is, the B level assignment of the shapes is not changed when testing for the main effect of A. Edgington and Onghena (2007) emphasize this point in permutation ANOVA tests for factorial experiments. Similarly, we next form several random arrangements of b groups (each with $N/b = an$ objects) to test for main effect B. Here we only permute between levels of factor B but not between levels of factor A. We are therefore assuming exchangeability of the observations \mathbf{X}_{ijl} among the factor levels when testing for the significance of the main effects.

To test for the significance of the main effects, for each random arrangement formed to test for factor A we compute the test statistic $F_0^{(1)}$. Likewise, for each random arrangement formed to test for factor B (notice two separate sets of random arrangements were formed, one set for each factor), we compute the test statistic $F_0^{(2)}$. The observed statistics F_0^{obs} for each test are then compared to the distribution of the corresponding F statistics obtained through random resampling. If r samples result in $F_0 > F_0^{obs}$, then the p -value of the test is given by r/P , where $P - 1$ denotes the number of random permutations generated for each effect tested, with the P th permutation being the actual observed arrangement, as discussed by Edgington and Onghena (2007).

Permutation tests for interactions in two-way ANOVA have been recently reviewed by Jung, Jhun, and Song (2006). These authors show that if random permutations of the observations are used to form ab groups of n objects from which (in our notation)

$$F_0^{(3)} = \left(n \sum_{i=1}^a \sum_{j=1}^b \frac{d_F^2(\bar{\mathbf{X}}_{ij.} - \bar{\mathbf{X}}_{i..} - \bar{\mathbf{X}}_{.j.} + \bar{\mathbf{X}}_{...}, \bar{\mathbf{X}}_{...})}{(a-1)(b-1)} \right) / \left(\sum_{i=1}^a \sum_{j=1}^b \sum_{l=1}^n \frac{d_F^2(\mathbf{X}_{ijl}, \bar{\mathbf{X}}_{ij.})}{ab(n-1)} \right) \quad (16)$$

is computed, the test for no AB interaction will have an inflated Type I error with respect to a Normal- F test, and it will have lower power to detect true differences. An alternative is to assume exchangeability across cells of the quantities \mathbf{X}_{ijl}^* defined as

$$\mathbf{X}_{ijl}^* = \mathbf{X}_{ijl}^p - \bar{\mathbf{X}}_{i..} - \bar{\mathbf{X}}_{.j.} + \bar{\mathbf{X}}_{...}$$

and then for each random permutation compute

$$F_0^{(3)*} = \left(n \sum_{i=1}^a \sum_{j=1}^b \frac{d_F^2(\bar{\mathbf{X}}_{ij.}^* - \bar{\mathbf{X}}_{i..}^* - \bar{\mathbf{X}}_{.j.}^* + \bar{\mathbf{X}}_{...}^*, \bar{\mathbf{X}}_{...}^*)}{(a-1)(b-1)} \right) / \left(\sum_{i=1}^a \sum_{j=1}^b \sum_{l=1}^n \frac{d_F^2(\mathbf{X}_{ijl}^*, \bar{\mathbf{X}}_{ij.}^*)}{ab(n-1)} \right).$$

According to Jung, Jhun, and Song (2006) this procedure has slightly lower Type I error than advertised (for the normal case), but the power is better than using (16), even for small sample sizes and nonnormal data. Note that before performing the permutations, all matrices \mathbf{X}_{ijl} need to be registered with respect to the overall mean using GPA.

5. A REAL MANUFACTURING APPLICATION

We now illustrate the aforementioned ANOVA tests with the titanium lathe-turning manufacturing example described in the [Introduction](#). In terms of the notation introduced earlier, we have $m = 2$, $a = b = 3$, $n = 10$, and $k = 64$. The two-way ANOVA test for shapes was applied to this experimental data. The results are shown in Table 2. The shape-space dimension is $M = (k-1)m - 1 - m(m-1)/2 = 124$, and hence, the degrees of freedom for A, B, AB, and error are $(a-1)M = (b-1)M = 248$, $(a-1)(b-1)M = 496$, and $ab(n-1)M = 10,044$, respectively.

Note how the sums of squares identity holds very exactly. This is because the shapes are very close to the average shape in shape space, a situation common in manufacturing data; in

Table 2. ANOVA for the shapes in the machining experiment

Source	$SS (\times 10^{-7})$	df	$MS (\times 10^{-10})$	F_0
A (depth)	9.88	248	39.85	7.98
B (speed)	4.85	248	19.56	3.91
AB int.	2.34	496	4.71	0.94
Error	50.14	10,044	4.99	
Total	67.22	11,036		




Figure 8. Scores of the first two principal components for all treatment groups in the titanium experiment. The data is concentrated and varies little, hence the shape space curvature can be neglected and the ANOVA identity holds very closely.

this example we have that $\max d_F = 5.13 \times 10^{-4}$. Figure 8 gives a plot of the principal component (PC) scores for each treatment combination over the first two PC's of matrix \mathbf{X}_v (see Section 2.2). This shows that the Gaussian assumption seems to hold, but reveals that the isotropy assumption is suspect. This plot also shows the data to be very concentrated around the mean shape.

Figure 9 shows the shape main effect plots and Figure 10 shows the shape interaction plots for this experiment. As can be seen, both depth (A) and speed (B) induce an oval (bilobed) deformation in the mean profile, but the induced ovals are oriented differently. If we apply the F distribution results suggested by Goodall (1991), the resulting p -values are 0.0000, 0.0000, and 0.8039, so we would conclude both depth and speed affect the mean shape of the parts but there is no depth-speed interaction.

For this example, the permutation ANOVA test, which given the evidence against isotropy should be the preferred test, indicates a similar conclusion to the F test. One thousand permutations ($P = 999$) were generated to test each of the three hypothesis. The empirical CDF of the tests is shown on Figure 11. The empirical p -values are 0.0, 0.0, and 0.5180, which are in general agreement with the ANOVA F shape test: the main effects of both factors are significant but the interaction is not. Finally, the p -values for the ANOVA on the size of the parts described in Section 3.3 yielded p -values of 0.7372, 0.6841, and 0.9647, thus there is no evidence that the factors (or their interaction) affect the size of the parts.

5.1 Shape Optimization

The shape effect plots presented above are useful tools to visualize the effect of each factor on the *mean* shape. In general, in advanced manufacturing data the mean shape will not be exactly equal to the nominal, or target, shape \mathbf{T} , but is typically quite close to it. To determine the optimal factor levels that will get the shape of the objects closest to \mathbf{T} , it is necessary to contrast the deviations between the cell means $\bar{\mathbf{X}}_{ij.}$ and the nominal shape \mathbf{T} . These differences can then be displayed graphically




Figure 9. Main effect on shapes for A = depth (a) and B = speed (b). The circles denote the overall mean shape $\bar{\mathbf{X}}_{...} = \hat{\boldsymbol{\mu}}$. The effects have been increased 1000 times to allow visualization.

as in Figure 12. The nine graphs in the figure correspond to the 3×3 cell means for the Titanium lathe experimental data. The arrows are proportional to the deviations from each cell mean to a perfect circle with 64 landmarks. To prepare these plots, the nominal and the cell means must have corresponding landmarks and be registered using the procrustes algorithm. In the example, it is obvious that factor B (speed) should not be set at its low level since this will result in large deviations from nominal. Likewise, setting factor A (depth of cut) to its low level is clearly best. The plots are of value since an engineer can visualize the directions and relative dimensions of the deviations at each landmark. A nongraphical measure useful to determine

optimal settings is the full procrustes distance $d_F^2(\bar{\mathbf{X}}_{ij}, \mathbf{T})$ computed for $i = 1, \dots, a$ and $j = 1, \dots, b$. We would choose the levels i, j that minimize this quantity. Applying this idea to the titanium lathe experimental data, the 3×3 matrix of distances $[d_F^2(\bar{\mathbf{X}}_{ij}, \mathbf{T})]$ is equal to

$$10^{-6} \times \begin{pmatrix} 0.299 & 0.266 & 0.282 \\ 0.354 & 0.340 & 0.291 \\ 0.556 & 0.409 & 0.334 \end{pmatrix}$$

and hence, it would appear there is some evidence that the best settings are depth = A = low and speed = B = medium or high. From a manufacturing point of view these settings make




Figure 10. AB shape interaction effect plot. The effect of depth (A) when changed from low to high while varying the speed (B) is similar to that observed in the main effect plot of depth (A) alone [Figure 9(a)], hence, there is no evidence of an AB interaction. The effects have been increased 1000 times to allow visualization.

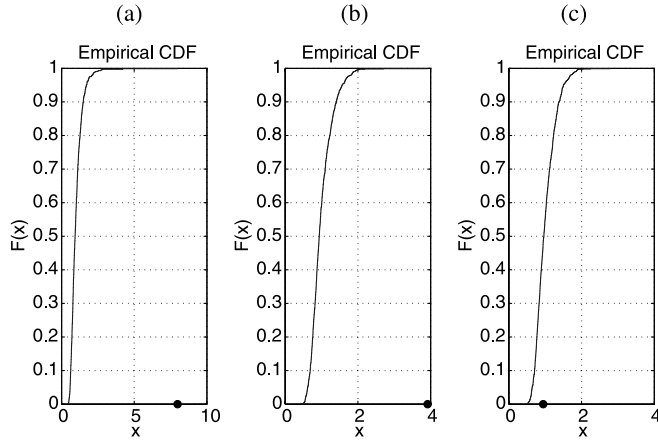


Figure 11. Empirical cumulative distribution functions for the permutation ANOVA shape test, 1000 permutations. (a) Main effect of depth-A (p -value is 0.0), (b) main effect of speed-B (p -value = 0.0), (c) AB interaction (p -value = 0.518). The dots on the x -axis indicate the observed test statistics, F_0^{obs} .

physical sense, since using the lowest depth of cut implies lowest cutting forces and this can help obtaining a perfect circle. Using a relatively higher speed can allow the turning spindle to encounter less rotation problems related to an oval contour shape, which is common in this data set (note how a high level for speed is best if the depth of cut is medium or high, and only second best if the depth is low).

Note, however, that the d_F^2 values are subject to sampling variability and hence the deviations from nominal at the different factor levels may not be statistically different. What we have in effect is a *multiple comparisons problem for shapes*: we wish to compare the deviations from nominal $\mu_{ij} - \mathbf{T}$ to choose the optimal factor settings. Since the nominal shape \mathbf{T} is

a constant, it suffices to make multiple comparisons on the cell means alone. A formal two-sample test for shape differences due to Goodall (1991) can be used for this purpose. Under the normal isotropic assumptions, the hypothesis $H_0: [\mu_{ij}] = [\mu_{i'j'}]$ is tested against a general alternative with the F statistic:

$$F_0 = \frac{n_1 + n_2 - 2}{1/n_1 + 1/n_2} \times \frac{G(\bar{\mathbf{X}}_{ij}, \bar{\mathbf{X}}_{i'j'})}{G(\mathbf{X}_{ij1}, \mathbf{X}_{ij2}, \dots, \mathbf{X}_{ijn_1}) + G(\mathbf{X}_{i'j'1}, \mathbf{X}_{i'j'2}, \dots, \mathbf{X}_{i'j'n_2})} \sim F_{M, (n_1 + n_2 - 2)M}. \quad (17)$$

The usual caveats as in any multiple comparisons problem apply. If a preset significance level is used, the overall (experimentwise) Type I error rate may be higher than advertised. A Bonferroni correction could be used to reduce this problem, given that no other multiple comparison procedure with control of the experimentwise Type I error rate exists for shape analysis (a permutation test based on F_0 above can be performed instead similarly to the permutation tests of Section 4, but these will have the same Bonferroni-like problems). In addition, the comparisons to make should be planned before observing the experimental results. To illustrate the multiple comparison tests, suppose before running the titanium lathe experiment we suspect (based on prior process knowledge) that the lowest value for depth of cut should improve circularity. Thus, we wish to make all possible multiple comparisons when factor A is set to its low setting. Table 3 shows the results from the corresponding tests. As it can be seen, the cell mean for the setting A = low, B = medium ($\mu_{1,2}$) is not significantly different than that for the setting A = low, B = high ($\mu_{1,3}$). Hence, either of these two settings seems to provide equally good circular shapes.

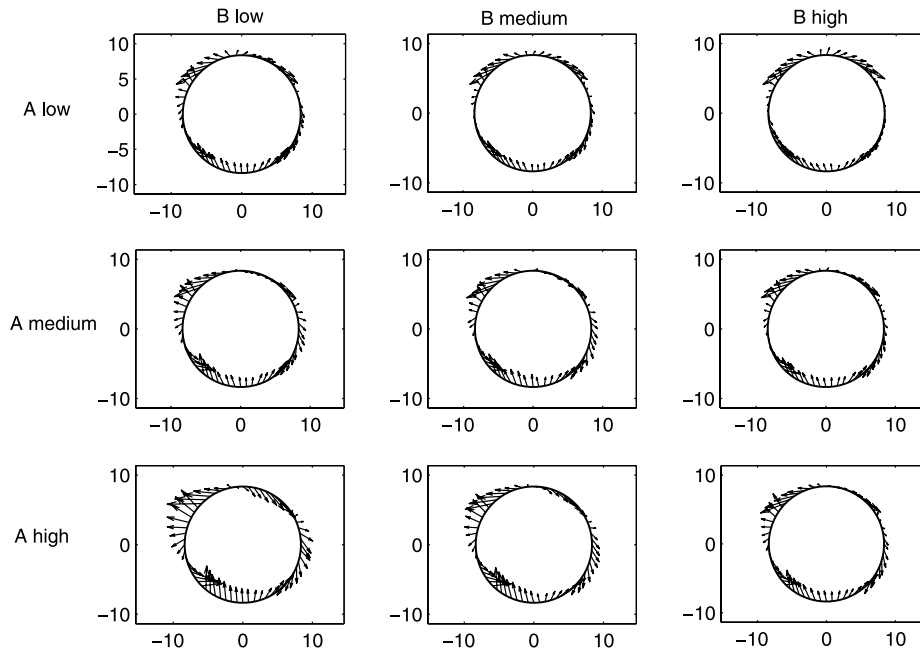


Figure 12. Deviations from nominal plots for the titanium lathe experiment. The arrows in each plot i, j are the deviations $\bar{\mathbf{X}}_{ij} - \mathbf{T}$ shown at each landmark in the target shape (a perfect circle) multiplied by a scale factor equal to 400 to allow visualization.

Table 3. Multiple comparisons for the cell means when factor A = low ($i = 1$), titanium lathe experiment

$\mu_{i,j}$		vs.	$\mu_{i',j'}$		p -value for $H_0: [\mu_{ij}] = [\mu_{i'j'}]$
i	j		i'	j'	
1	1		1	2	0.0049
1	1		1	3	0.0000
1	2		1	3	0.6893

6. POWER ANALYSIS OF THE DIFFERENT METHODS FOR CIRCULAR SHAPE DATA

We study the statistical performance of the following three methods used to determine differences in circular shapes: (1) the ANOVA test based on Goodall's F test (Section 3), (2) the permutation test of Section 4, and (3) an ANOVA performed on the *circularity form error*, the usual way to determine factor effects on the circularity of manufactured parts (Krulikowski 1996; ASME 1994). The circularity (roundness) form error is defined in manufacturing practice as the smallest difference between the radii of two coaxial circles that contain all the measurements. Note that the computation of the circularity form error does not require labeled (corresponding) landmarks. A standard ANOVA was then conducted on these form errors. Circular shape profile data was simulated that is less circular as the level of a factor increases (we only consider the one-factor ANOVA case in this section). As the value of the factor—say, depth of cut—changes from low to high, a second harmonic with amplitude δ was added to the simulated circular profiles. To determine a power curve for each ANOVA method, different values of the “shift” in amplitude δ were tried. Since the ability to detect a change in a circular profile is related to both the amount of noise *and* the magnitude of the radius R of the ideal circular profile, the number δ was expressed as a multiple w of the ratio σ/R , thus the shift in amplitude was given by $\delta = w\sigma/R$. The values tested for w were $\{0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$, with $w = 0$ ($\delta = 0$) denoting the baseline case where the factor has no effect on the circular contour. This double standardization of the shift in amplitude of the second harmonic allows us to observe the behavior of the different methods for a circle of any size since the graph gives power as a function of w . That is, as long as the added amplitude is expressed in multiples of σ/R , the power curves for the different methods will be as shown in Figure 13.

Twenty parts ($n = 20$) were simulated at each of the two levels of the factor (the low value always giving a circular contour and the high value always adding a second order bilobe of amplitude δ). For each value of the w multiplier (the “shift size”), 100 simulations were performed, and $N(0, \sigma^2)$ iid errors were added at each landmark, where $k = 64$ equidistant landmarks on the circle were used. The percentage of times p that the null hypothesis of no factor effect was rejected is the estimated power of each test. The desired probability of a Type I error was set at 5%. A value of $\sigma/R = 0.01$ with $\sigma = 0.05$ and $R = 5$ was used to generate the plot.

From Figure 13 it is clear that the ANOVA on the circularity form error is unable to detect differences in the circular profiles unless these become very large. This relatively low power

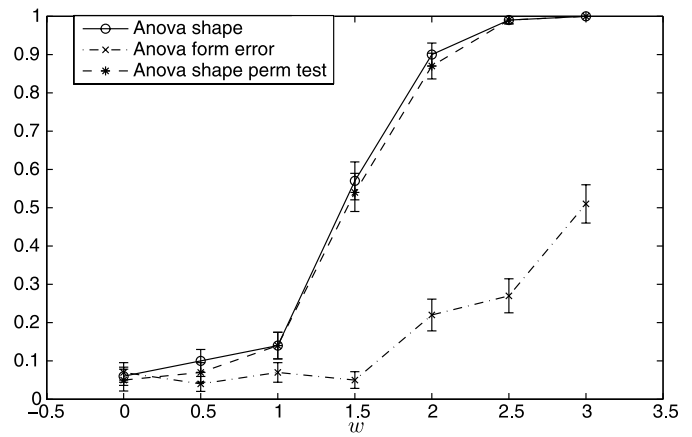


Figure 13. Power results for circular profiles ($\alpha = 0.05$) for the one-factor ANOVA based on the circularity form error, the ANOVA F -test for shapes, and the permutation ANOVA test for shapes (100 permutations). The estimated power is valid for any circular objects as long as the amplitude of the added second harmonic is $w\sigma/R$ with w as in the x -axis. Bars indicate 95% confidence intervals.

is probably the result of using only two landmarks, the closest and the farthest from the center, when computing the circularity form error, as opposed to using all the landmarks as in the proposed SSA methods. Given that the isotropic error assumptions are frequently unrealistic, we suggest the *permutation* ANOVA on the shapes as the test to use in practice. The test achieves the desired α level and has high probability of detection, almost as high as for the F -test (under ideal conditions for the latter) and has much higher power than the traditional form error method used in manufacturing practice. Furthermore, the permutation test does not rely on the normality or anisotropic variance assumptions. This test, used in combination with the visualization techniques described in Section 3 provides a powerful set of tools to analyze experiments when the responses of interest are the geometrical shapes of objects.

7. DISCUSSION AND FURTHER WORK

In this article, new two-way ANOVA tests for shapes have been presented. The power of the tests to detect effects in circular mean shapes was investigated and new graphical tools were presented, namely, shape main effect plots, shape interaction plots, and deviation from nominal shape plots. These graphics aid in the analysis and interpretation of experiments where the response of interest is the shape of manufactured objects.

We suggest using the permutation ANOVA test when analyzing shape responses in a designed experiment. The test achieves the desired α level and has high power, almost as high as for the ANOVA F -test for shapes and much higher than the usual form error ANOVA performed in manufacturing practice. Furthermore, it does not rely on the normality or isotropic variance assumptions as the F test does. Further more general evidence supporting this recommendation can be found in Alshraideh and Del Castillo (2009), who conducted a power analysis for all the tests for shape difference presented here for a variety of shapes, both of interest in manufacturing and for arbitrary shapes. They also consider the nonisotropic variance case, since

in the present article we only have considered the isotropic variance case.

For the type of manufacturing applications we study, the shapes are very close to the overall mean and this satisfies an underlying assumption of the ANOVA methods presented. When shapes differ considerably, the nonlinearity of the shape space cannot be neglected, and the ANOVA identity will not hold. Recent work along the lines in Huckemann, Hotz, and Munk (2010a, 2010b) should be applied instead. These authors present shape analysis methods on the shape space. A difference with the type of applications in the SSA literature (including the recent work by Huckemann, Hotz, and Munk) is that in manufacturing, the number of landmarks k is usually in the hundreds, whereas in zoology and botany applications k is frequently less than 10.

There are several lines of further work that can be undertaken. A basic problem is to *label* the points in each part such that they relate to similar physical locations (landmarks) between parts, that is, the points in different parts are corresponding. This is a problem that has received considerable attention in the field of computer vision (see, e.g., Belongie, Malik, and Puzicha 2002; Gold et al. 1998; and for a more recent presentation and review, Davies, Twining, and Taylor 2008) where it is called the “matching points” or the “correspondence” problem.

One immediate application of the graphical techniques presented in this article is related to the “Robust Parameter Design” (RPD) problem (Montgomery 2006). In an RPD problem, experimental factors are divided in two categories, controllable (C) and noise (N), and the purpose of the RPD experiment is to find the values of the controllable factors that will yield response values that are robust (i.e., not sensitive) with respect to random variation in the noise factors. The easiest way to achieve this goal in a standard factorial experiment is to look at the $C \times N$ interaction plots, and choose values of the controllable factors where the slope of the interaction plot is smallest (this minimizes the transmitted noise). We could use either the AB interaction plot or the deviation from nominal plots in an analogous way. For example, in Figure 10 (or Figure 12), suppose A is controllable but B is a noise factor. Then the high value of A should not be selected since this results in shapes that are most sensitive to variations in factor B.

Another area of potential further work, suggested to us by an associate editor, is to modify the ANOVA tests presented for unbalanced experiments. Although the power results presented in Section 6 refer to balanced experiments, the one-way ANOVA test is valid for general unbalanced designs. Similarly, the test for comparing means used to determine best factor combinations [Equation (17)] is valid for unequal sample sizes. An easy unbalanced case to handle (in any number of factors) is that of a few missing observations, which could occur in manufacturing perhaps due to defective CMM data acquisition on a given part. While this can always be treated by simply remeasuring the part, this can also be treated by replacing the missing observation with the mean shape for the remaining shapes in the cell, $\hat{\mu}_{i,j} = \bar{\mathbf{X}}_{i,j}$. Fortunately, in experimental designs in manufacturing this is the most common type of unbalancing found.

Another consideration suggested by the associate editor is that of shapes serially correlated with each other (between shape correlation), rather than shapes with a (nonisotropic) covariance structure within each shape (between landmarks). This

in principle could occur in manufacturing due to wearing-out phenomena, although for the typical small sample sizes in shape analysis experiments the wear would have to be quite drastic to be observable, so we conjecture that this is a rare situation in practice. Serial correlation violates an underlying assumption not only in the ANOVA shape methods presented here but also in “standard” ANOVA. The suggested ANOVA permutation test is particularly sensitive to serial correlation since it is based on an assumption of exchangeability. Just as in standard situations, if the run order cannot be randomized (otherwise the permutation test is still valid), one is forced to either model the correlation somehow (e.g., via time series analysis techniques) or to anticipate it in the experimental design (e.g., via blocking in the case of time trends). How to do this in experiments with shape responses is a matter of further work.

This article has focused on 2D shapes. Extensions to the three-dimensional case are evidently practical (see Alshraideh and Del Castillo 2009 for some recent results on the ANOVA tests applied to 3D objects). We have also confined ourselves to least squares registration methods (i.e., procrustes), but the use of registration methods that are more robust to wild landmark locations (e.g., Rohlf and Slice 1990) can be investigated. A comparative study of the statistical power that the (permutation) ANOVAs have when based on different registration methods also seems necessary.

Finally, this article has focused on shape analysis methods based on procrustes methods. Comparisons with shape analysis methods based on the Euclidean distances between matrices (e.g., Lele and Richtsmeier 2001) is also a matter for further research. Alshraideh and Del Castillo (2009) show how distance-based methods seem to have low power in the case of 2D circles and 3D cylinders, but do have relatively competitive power compared to the ANOVA methods presented here to detect differences in certain geometrical objects that have just a few landmarks.

APPENDIX: COMPUTER IMPLEMENTATION OF THE ANOVA SHAPE ALGORITHMS

Matlab programs that perform the computations required for ANOVA for shapes, including visualization, were written for this research and can be downloaded from <http://www2.ie.psu.edu/Castillo/research/EngineeringStatistics/software.htm>.

There are three main programs: one for the two-way ANOVA F -test for shapes that assumes normality and isotropy and produces the plots helpful for visualization (ANOVAShape.m); a second program that implements the (recommended) two-way permutation ANOVA test for shapes, valid under more general error distribution assumptions (permutationANOVAShape.m) and a third program, ANOVAShapeNominal.m, that computes and displays the deviation from nominal plots and performs multiple comparisons of cell means. The dataset for the lathe experiment used in this article can also be downloaded from the same URL.

ACKNOWLEDGEMENT

We wish to thank the Editor, an Associate Editor, and two referees for several helpful comments and suggestions that allow us to improve and clarify the presentation of this work.

Thanks to Mr. Hussam Alshraideh (PSU) for his help with Figure 13. This research was funded in part by the National Science Foundation grant CMI-0825786 and by the Italian Ministry of Education, University and Research (MIUR) grant FIRB RBIP069S2T 005.

[Received November 2008. Revised August 2010.]

REFERENCES

- Adams, D. C., Rohlf, F. J., and Slice, D. E. (2004), "Geometric Morphometrics: Ten Years of Progress Following the "Revolution"," *Italian Journal of Zoology*, 71, 5–16. [2,3,6]
- Alshraideh, H., and Del Castillo, E. (2009), "Statistical Performance of Tests for Mean Shape Difference With Application in Manufacturing," working paper, Engineering Statistics Laboratory. [13,14]
- American Society of Mechanical Engineers (ASME) (1994), *ASME Y14.5M-1994, Dimensioning and Tolerancing*, New York: ASME. [2,3,13]
- Belongie, S., Malik, J., and Puzicha, J. (2002), "Shape Matching and Object Recognition Using Shape Contexts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24 (24), 509–522. [14]
- Bookstein, F. L. (1986), "Size and Shape Spaces for Landmark Data in Two Dimensions," *Statistical Science*, 1, 181–242. [1]
- Colosimo, B. M., Pacella, M., and Semeraro, Q. (2008), "Statistical Process Control for Geometric Specifications: On the Monitoring of Roundness Profiles," *Journal of Quality Technology*, 40 (1), 1–18. [2]
- Davies, R., Twining, C., and Taylor, C. (2008), *Statistical Models of Shape, Optimisation and Evaluation*, London: Springer-Verlag. [3,14]
- Dryden, I. L., and Mardia, K. V. (1998), *Statistical Shape Analysis*, Chichester: Wiley. [2,3,5-7,9]
- Edgington, E. S., and Onghena, P. (2007), *Randomization Tests* (4th ed.), Boca Raton, FL: Chapman & Hall/CRC Press. [9]
- Gold, S., Rangarajan, A., Lu, C.-P., Pappu, S., and Mjolsness, E. (1998), "New Algorithms for 2D and 3D Point Matching: Pose Estimation and Correspondence," *Pattern Recognition*, 31 (8), 1019–1031. [14]
- Goodall, C. (1991), "Procrustes Methods in the Statistical Analysis of Shape," *Journal of the Royal Statistical Society, Ser. B*, 53 (2), 285–339. [3,5-7,10,12]
- Gower, J. C. (1975), "Generalized Procrustes Analysis," *Psychometrika*, 40 (1), 33–51. [5]
- Horn, B. K. P., Hilden, H. M., and Negahdaripour, S. (1988), "Closed-Form Solution of Absolute Orientation Using Orthonormal Matrices," *Journal of the Optical Society of America A*, 5, 1127–1135. [5]
- Huckemann, S., Hotz, T., and Munk, A. (2010a), "Intrinsic MANOVA for Riemannian Manifolds With an Application to Kendall's Space of Planar Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32 (4), 593–603. [6,14]
- (2010b), "Intrinsic Shape Analysis: Geodesic PCA for Riemannian Manifolds Modulo Isometric Lie Group Actions," *Statistica Sinica*, Preprint SS-07-173. [6,14]
- Jackson, J. E. (2003), *A User's Guide to Principal Components*, New York: Wiley. [5]
- Jung, B. C., Jhun, M., and Song, S. H. (2006), "A New Random Permutation Test in ANOVA Models," *Statistical Papers*, 48, 47–62. [10]
- Kang, L., and Albin, S. L. (2000), "On-Line Monitoring When the Process Yields a Linear Profile," *Journal of Quality Technology*, 32, 418–426. [2]
- Kendall, D. G. (1984), "Shape Manifolds, Procrustean Metrics, and Complex Projective Spaces," *Bulletin of the London Mathematical Society*, 16, 81–121. [1,3,4]
- (1989), "A Survey of the Statistical Theory of Shape," *Statistical Science*, 4 (2), 87–89. [3]
- Kendall, D. G., Barden, D., Carne, T. K., and Le, H. (1999), *Shape and Shape Theory*, New York: Wiley. [4]
- Klein, F. (1939), *Elementary Mathematics From an Advanced Standpoint: Geometry*, New York: The Macmillan Company. [1]
- Klingenberg, C. P., and McIntyre, G. S. (1998), "Geometric Morphometrics of Developmental Instability: Analyzing Patterns of Fluctuating Asymmetry With Procrustes Methods," *Evolution*, 52 (5), 1363–1375. [6,9]
- Klingenberg, C. P., and Monteiro, L. R. (2005), "Distances and Directions in Multidimensional Shape Spaces: Implications for Morphometric Applications," *Systematic Biology*, 54 (4), 678–688. [3,6]
- Krim, H., and Yezzi, A., Jr. (eds.) (2006), *Statistics and Analysis of Shapes*, Boston: Birkhäuser. [3]
- Krulikowski, A. (1996), *Geometric Dimensioning & Tolerancing (A Companion to the ASME Y14.5M-1994 Dimensioning & Tolerancing Standard)*, Westland, MI: Effective Training Inc. [2,13]
- Langron, S. P., and Collins, A. J. (1985), "Perturbation Theory for Generalized Procrustes Analysis," *Journal of the Royal Statistical Society, Ser. B*, 47 (2), 277–284. [6]
- Lele, S. R., and Richtsmeier (2001), *An Invariant Approach to Statistical Analysis of Shapes*, Boca Raton, FL: Chapman & Hall/CRC. [14]
- Mardia, K. V. (1989), "Comment: Some Contributions to Shape Analysis," *Statistical Science*, 4 (2) 108–111. [3]
- Miller, A., and Wu, C. F. J. (1996), "Parameter Design for Signal-Response Systems: A Different Look at Taguchi's Dynamic Parameter Design," *Statistical Science*, 11, 122–136. [2]
- Montgomery, D. C. (2006), *Design and Analysis of Experiments* (6th ed.), New York: Wiley. [8,14]
- Nair, V. N., Taam, W., and Ye, K. Q. (2002), "Analysis of Functional Responses From Robust Design Studies," *Journal of Quality Technology*, 34 (4), 355–370. [2]
- Press, S. J. (2005), *Applied Multivariate Analysis* (2nd ed.), Mineola, NY: Dover. [7]
- Rohlf, F. J., and Slice, D. (1990), "Extension of the Procrustes Method for the Optimal Superimposition of Landmarks," *Systematic Zoology*, 39 (1), 40–59. [14]
- Searle, S. R. (1971), *Linear Models*, New York: Wiley. [7]
- Snee, R. D. (1972), "A Useful Method for Conducting Carrot Shape Studies," *Journal of Horticultural Science*, 47, 267–277. [2]
- Snee, R. D., and Andrews, H. P. (1971), "Statistical Design and Analysis of Shape Studies," *Applied Statistics*, 20 (3), 250–258. [2]
- Srivastava, A., Damon, J. N., Dryden, I. L., and Jermyn, I. H. (2010), "Guest Editors' Introduction to the Special Section on Shape Analysis and Its Application in Image Understanding," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32 (4), 577–578. [3]
- Ten Berge, J. M. F. (1977), "Orthogonal Procrustes Rotation for Two or More Matrices," *Psychometrika*, 42 (2), 267–276. [5]