ORIGINAL ARTICLE

Functional data analysis and visualisation of three-dimensional surface shape

Stanislav Katina[1,2] | Liberty Vittert[3] | Adrian W. Bowman[4]

[1]Institute of Mathematics & Statistics, Masaryk University, Brno, Czech Republic

[2]Institute of Computer Science of the Czech Academy of Sciences, Prague, Czech Republic

[3]Olin Business School, Washington University in St. Louis, St. Louis, MO, USA

[4]School of Mathematics & Statistics, The University of Glasgow, Glasgow, UK

Correspondence

Adrian W. Bowman, School of Mathematics & Statistics, The University of Glasgow, Glasgow, UK.
Email: adrian.bowman@glasgow.ac.uk

Abstract

The advent of high-resolution imaging has made data on surface shape widespread. Methods for the analysis of shape based on landmarks are well established but high-resolution data require a functional approach. The starting point is a systematic and consistent description of each surface shape and a method for creating this is described. Three innovative forms of analysis are then introduced. The first uses surface integration to address issues of registration, principal component analysis and the measurement of asymmetry, all in functional form. Computational issues are handled through discrete approximations to integrals, based in this case on appropriate surface area weighted sums. The second innovation is to focus on sub-spaces where interesting behaviour such as group differences are exhibited, rather than on individual principal components. The third innovation concerns the comparison of individual shapes with a relevant control set, where the concept of a normal range is extended to the highly multivariate setting of surface shape. This has particularly strong applications to medical contexts where the assessment of individual patients is very important. All of these ideas are developed and illustrated in the important context of human facial shape, with a strong emphasis on the effective visual communication of effects of interest.

1 | INTRODUCTION

Statistical shape analysis is a research topic which has seen very substantial growth and development in recent years. Early work in this area focused on representations of shape through carefully chosen landmarks, as point locations with an interpretation which corresponds across different shapes. Dryden and Mardia (2016) provide a very comprehensive description of methods for the analysis of landmarks, but the later chapters of the book also indicate the much wider array of data types which are becoming available, driven by rapid advances in imaging technology. A particular example is the increasing availability of sensors which employ techniques such as laser scanning or stereo-photogrammetry to create high-resolution data on surface shape in three dimensions. This has a very wide variety of applications and it is the focus of the present paper. Figure 1 shows an image of a human face as an example of the kind of three-dimensional (3D) surface data which is now easily obtainable.

Single instances of 3D surface data can be displayed in a variety of ways; in particular, the `rgl` package (Adler and Murdoch, 2019) is an indispensable tool for those from the R (R Core Team, 2019) community, as it provides access to the *OpenGL* industry-standard tools for 3D display. However, the effective display of patterns and variation in collections of 3D objects is more challenging. Bowman and Bock (2006) gave some discussion of this for 3D points and curves, but the aim of the present paper is to provide new tools for the modeling and visualisation of samples of 3D surface data.

The starting point is a description of an individual surface which has a consistent meaning across all the surfaces in the data set. This can be approached in different ways and the particular method adopted here is described in Section 2. Some obvious issues of analysis then commonly arise. These include the need for methods to:

- register the surfaces to a common co-ordinate system;
- characterise the variation present in a sample of surfaces;
- compare surface shapes across groups;
- assess the surface shape of an individual against a relevant control set.

These problems are tackled here from a functional perspective. Adaptations of standard methods of Procrustes analysis are introduced in Section 2, using a metric based on an approximated surface
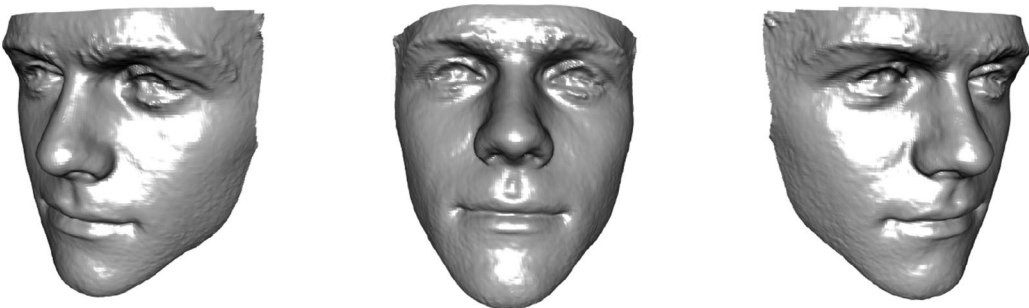

FIGURE 1   An example of a 3D facial image, at different orientations

integral rather than individual point locations. Non-linear registration through warping is also described as a means of displaying the results of analysis at higher resolution, for visual effect. Basic methods of visualising surface differences are also reviewed. Section 3 discusses the use of principal components in exploring the variation in surface data and in comparing groups. Again a functional perspective is adopted, based on surface integration. This section also discusses how effects can be visualised by characterising the shape changes associated with appropriate subspaces, rather than through examination of individual components. Section 4 addresses the situation where there is a need to assess the characteristics of individual surfaces, and in particular of any shape features which are not consistent with control shape. Some final discussion and reflection is provided in Section 5.

The methods proposed in the paper are illustrated throughout on images of human facial shape. There is a strong emphasis on the creation of visual displays which communicate patterns in the data, the evidence and nature of group differences, and the distinctive characteristics of individuals, as clearly as possible. Graphics are provided in static form but animations are also available in the *Supplementary Information*.

2 | SOME FUNDAMENTAL TOOLS FOR SURFACES

2.1 | Facial models

A model for an individual surface should provide a structured representation of shape whose components have a consistent interpretation across the other surfaces in the data set. This then allows the investigation of pattern and variation in shape. Landmarks satisfy this criterion and so, while the information they carry is limited, they have often been used as the starting point for more complex models. (Hammond et al, 2004; Mao et al, 2006; Paulsen and Hilger, 2003) give examples of this approach where a template of a human face is 'warped' onto an observed image. Landmarks on the template are transformed in a non-linear manner to match those on the image exactly, with the surface of the template then adjusted further to improve the match with the surface of the image. This might be done by locating closest points or by matching the characteristics of local surface curvature. The resulting transformed template then provides a model for the surface whose meaning corresponds across all the images in the data set.

In an alternative approach, Vittert, Bowman and Katina (2019) took the view that ridge and valley curves provide the key information on shape, as these capture the locations where curvature is strongest. The two left-hand images in Figure 2 give examples of facial curvature, here in the form of *Gaussian curvature* and *shape index*; see Appendix A for details. Curvature information can then be used to fit a model consisting of a set of ridge, valley or other geodesic curves, with landmarks as end-points. A full surface representation can easily be constructed by interpolation across the surface patches bounded by these curves, although Vittert, Bowman and Katina (2019) give examples where the focused representation based on curves alone can be more informative. An example of the resulting facial model is shown in the right-hand panel of Figure 2, which uses colour and size to indicate the hierarchical nature of the information captured in landmarks, curves and surface patches.

2.2 | Registration

A key issue in the analysis of shape is that the observed images do not necessarily lie in a common co-ordinate system. The process of data capture does not usually give each image the same origin or
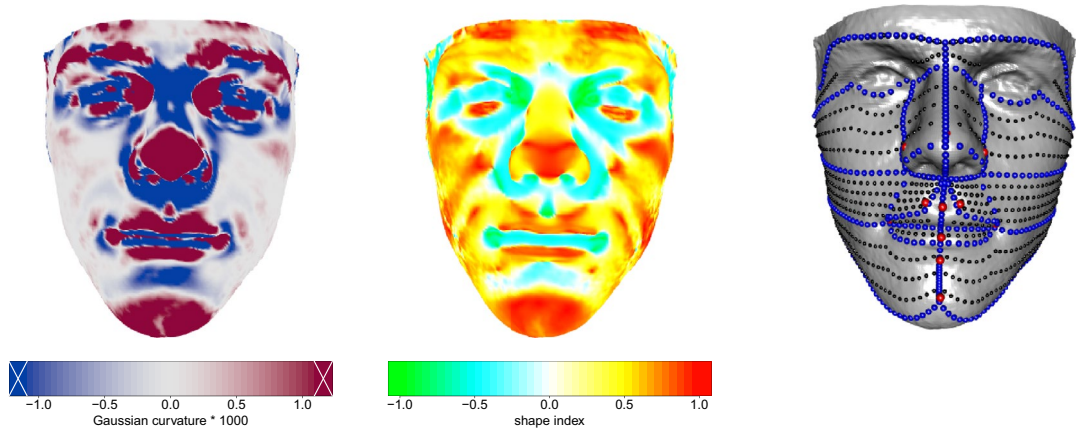
FIGURE 2   The two left-hand images show plots of Gaussian curvature and shape index. The crosses at the end of the Gaussian curvature colour scale indicate that exceptionally high values have been truncated. The right-hand image shows a fitted facial model, with landmarks (red), anatomical curves (blue) and surface patches (black) superimposed on an observed image (grey) [Colour figure can be viewed at wileyonlinelibrary.com]

orientation. The relative sizes of the images may also be viewed as unimportant from a shape perspective. It is therefore necessary to remove these extraneous aspects before statistical analysis. Different approaches to this are outlined below. These are described in the context of transforming an image X to match a reference image Y, where X and Y are $J \times 3$ matrices whose rows give the 3D positions of a fitted model in the discretised form of J point locations. The process of transforming X to match Y is referred to as registration.

2.2.1 | Procrustes methods

A very effective approach is to find the rotation matrix Γ, scaling parameter β and translation vector γ which bring X as close to Y as possible. Adopting a similar notation to Dryden and Mardia (2016, chapter 7), the method is expressed as:

$$\min_{\beta,\Gamma,\gamma} ||Y - \beta X\Gamma - 1_J\gamma^T||^2 = \min_{\beta,\Gamma,\gamma} \sum_{j=1}^{J} ||y_j - \beta\Gamma^T x_j - \gamma||^2, \tag{1}$$

where $||Z|| = \sqrt{\text{trace}(Z^TZ)}$ is the Euclidean norm, X and Y are assumed to have centroid 0, the column vectors x_j, y_j correspond to the jth rows of X and Y, and 1_J denotes a column vector containing J 1's. This is a standard example of an approach referred to as Procrustes registration. The ideas and methods of implementation involved are comprehensively described by Dryden and Mardia (2016, chapter 7).

These methods arose in the context of shape representations based on landmarks, often well-separated spatially. The representations we are now dealing with may have a discrete point-based form, for convenience, but they represent a continuous surface. This leads immediately to a functional data analysis perspective, as described by Ramsay and Silverman (1997). In the current setting, functional registration is achieved through

$$\min_{\beta,\Gamma,\gamma} \int_{S_Y} ||y - \beta\Gamma^T x(y) - \gamma||^2 dy, \tag{2}$$

where S_Y denotes the surface indexed by Y. The function $x(y)$ indexes the point x on the surface S_X which has geometrical correspondence with the point y on surface S_Y. This correspondence is established by the model, which uses the same geometrical features to characterise each surface and so, for example, the set of points which provide a discrete representation of each nose ridge have a point-by-point correspondence. The integral can now be approximated in discrete form as

$$\sum_{j=1}^{J} a_j ||y_j - \beta \Gamma^T x_j - \gamma||^2 = ||\sqrt{A}Y - \beta \sqrt{A}X\Gamma - \sqrt{A}1_J \gamma^T||^2, \tag{3}$$

where the weight a_j gives the surface area which surrounds point y_j and A is a diagonal matrix containing the a_j.

The weights a_j can be calculated easily from a surface triangulation associated with the discrete representation of each model. If $\mathcal{T}_1, \ldots, \mathcal{T}_T$ denote the set of surface triangles and $\mathcal{N}_j$ is the set of indices of triangles which have x_j as a vertex and $|\cdot|$ denotes area, then the weights are simply $a_j = \frac{1}{3}\sum_{t \in \mathcal{N}_j} |\mathcal{T}_t|$. The divisor 3 apportions one third of the area of each triangle to each of its three vertices. Clearly the approximation accuracy of the integral will increase with the density of the discrete representation. More sophisticated numerical methods are available; see Reeger, Fornberg and Watts (2016), for example.

The effect of the weighted registration (3) is substantially to remove the effect of the particular form of discrete representation used, together with its associated triangulation, by approximating the surface integral. In contrast, the unweighted version (1) has the undesirable property that it will give greater weight to those locations where the discrete representation happens to be more dense. It is not possible to create a discrete representation with uniform spacing across an arbitrary manifold.

Expansion of the right-hand side of expression (3), following the derivation of the unweighted case in Dryden and Mardia (2016, chapter 7), shows that the minimum is achieved when the matrices AX and AY are column-centred, with

$$\hat{\gamma} = 0,$$
$$\hat{\Gamma} = UV^T,$$
$$\hat{\beta} = \text{tr}\left\{Y^T AX\hat{\Gamma}\right\} / \text{tr}\left\{X^T AX\right\},$$

where $Y^T AX = ||\sqrt{A}Y|| \, ||\sqrt{A}X|| \, V\Lambda U^T$, with Λ diagonal. Brignell, Dryden and Browne (2015) discuss more complex forms of weighting for other purposes.

The case of matching one shape X to another Y is referred to as *ordinary Procrustes registration*. This provides the building block for *generalised Procrustes registration* which seeks a common registration of multiple shapes $X_1, \ldots, X_n$. The aim now is to minimise the sum of the deviations of transformed shapes from a common mean surface S_M with discrete model representation M. This can be expressed in functional form as

$$\sum_{i=1}^{n} \int_{S_M} ||m - \beta_i \Gamma_i^T x_i(m) - \gamma_i||^2 dm$$
$$\approx \sum_{i=1}^{n} ||\sqrt{A}M - \beta_i \sqrt{A}X_i \Gamma_i - \sqrt{A}1_J \gamma_i^T||^2,$$

with the transformation parameters $\beta_i, \Gamma_i, \gamma_i; i = 1, \ldots, n$. The weights in the diagonal matrix A are now the areas surrounding the vertices of the mean shape M. Following again the general structure outlined in Dryden and Mardia (2016, chapter 7), and beginning with each AX_i column centred, minimisation can be achieved by successive ordinary weighted Procrustes registration of the adjusted shapes

$X_i^P = \beta_i X \Gamma_i + 1_J \gamma_i^T$ onto the mean M, which itself is estimated simply as the average of the X_i^P. A size constraint is required to ensure that the solution does not degenerate to 0. From a functional perspective, size is expressed in the surface area of each shape and this is easily calculated in discrete form as the trace of A. Notice that integration is carried out over the mean surface S_M, so that A changes with each iteration.

Figure 3 provides an illustration from a sample of 61 males and 69 females, all adults of UK origin, where Procrustes registration has been applied to each sex separately and the resulting means matched by a further Procrustes step. Human sexual dimorphism has been extensively studied (see, e.g. Armann and Balthoff, 2012; Bruce et al, 1993; Claes et al, 2012; Wilkinson, 2004). Figure 3 highlights the key differences on average, with males exhibiting greater prominence in nasal, chin and brow ridge areas while females correspondingly exhibit more prominent cheeks.

This simple example also raises a key issue of visualisation, namely how to compare two 3D surfaces most effectively. The Figure adopts the simple strategy of superimposing the surfaces which gives a clear indication of which of the two shapes is more prominent in each area. Other strategies will be considered in Section 3. Colour choice is also an important aspect of visualisation and this is discussed very helpfully by Zeileis, Hornik and Murrell (2009), with effective solutions for comparing groups and displaying values on a continuous scale, including the presence of a reference value. These colour choices have been adopted through the paper.

In order to avoid noisy areas of the image around the eyes and nostrils, the facial models displayed in Figure 3 reduce the description at the brow and columella (between the nostrils) to two curves. These curves are included in surface-based registration, and in later analysis, simply by considering a small patch around each curve point, with area set to the average of the areas surrounding the surface points.

2.2.2 | Warping

Procrustes registration brings the co-ordinates of shape X as close as possible to those of shape Y, within the limits imposed by the use of translation, rotation and scaling only. However, there are situations where it is useful to match these two shapes exactly. This arises in some of the methods involved in constructing facial models, described in Section 2.1, where a template is initialised on an observed image by exact matching of a set of landmarks. A further example is in the improvement of the visual comparison of shapes such as those in Figure 3. If a high-resolution facial template is available, with an embedded shape model Z which corresponds to that of X and Y, then a smooth function which transforms Z to X exactly can be identified, in a process known as *warping*. This function can then be applied to the template to create a visual display which has smoother and more attractive surfaces than the lower resolution model X and which adds in detailed features such as nostrils and eyes, giving a
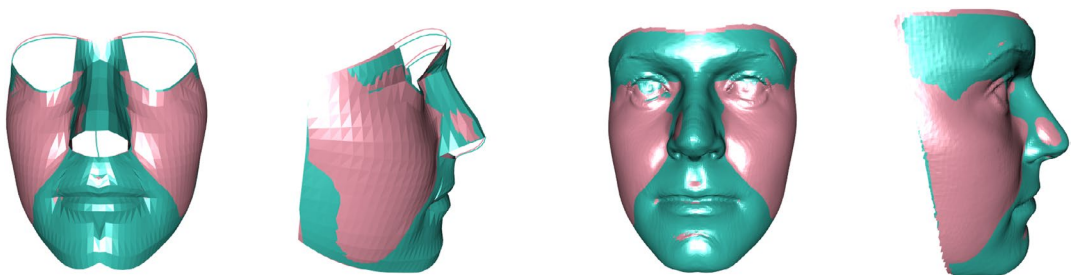

FIGURE 3   Procrustes matched male (green) and female (pink) mean shapes, in frontal and lateral orientations, using model meshes (left two images) and warped facial templates matched to these mean shapes (right two images). [Colour figure can be viewed at wileyonlinelibrary.com]

more effective and interpretable display of a human face. While care should be taken not to interpret the form of these very detailed features, the principal characteristics of the display all reflect the underlying model. Use of a template can also help to anonymise individual faces. The right-hand images in Figure 3 show the effects of employing templates in this way on both male and female means.

In the analysis of 2D shapes based on landmarks, the concept of a deformation grid to describe shape change is a very old one; see Thompson (1917, chapter 9). This uses a function which maps one set of landmarks to another exactly but, as it is expressed in functional form, this function can also be applied to a regular grid of locations over the first image to create a warped grid which expresses the underlying transformation. Methods based on pairs of thin-plate splines were first introduced in 2D by Bookstein (1989) and developed further by Bookstein (1997, chapter 7). The topic is also explained clearly by Dryden and Mardia (2016, chapter 12). Corresponding methods in 3D were first introduced by Gunz, Mitteroecker and Bookstein (2005) and applied to skulls by Mitteroecker et al. (2004) and Mitteroecker and Bookstein (2008), to long bones by Frelat et al. (2012), and to mice heads by Waddington et al. (2017). The literature on radial basis functions uses the same techniques but employs a different language. The technical details of warping in 3D are described in Appendix B.

2.3 | Visual comparison of two shapes

The male and female example of Figure 3 raises the question of how two shapes can most effectively be compared visually. The challenge is that in addition to the 3D shapes themselves, comparison involves a vector field of differences, with a displacement vector at each position on the individual shapes. A helpful strategy is to display one shape and use colour to inform on the shape difference from the other shape at each location. Figure 4 illustrates this by plotting the female mean face and using colour to indicate distances to the male mean. The distances in the separate x, y and z co-ordinates are shown. These co-ordinates can be given clear interpretations by orienting the female mean so that nominated landmarks such as the outer corners of the eyes (*exocanthions*) define the direction of the x-axis and others such as the top of the nose ridge between the eyes and the central point at the base of the nose (*nasion* and *subnasale*) define the y-axis. The final two images use colour to indicate the distance between the corresponding points on the female and male means projected along the normal direction at the surface of the female mean, and the Euclidean distance between corresponding points, with sign determined by whether the projection along the normal is positive or negative.

None of these devices captures the information in the shape difference completely, because the change is in 3D while the colour scale can represent only a single dimension, but they provide options for detailed exploration. The simple superimposition of surfaces illustrated in Figure 3 is usually a
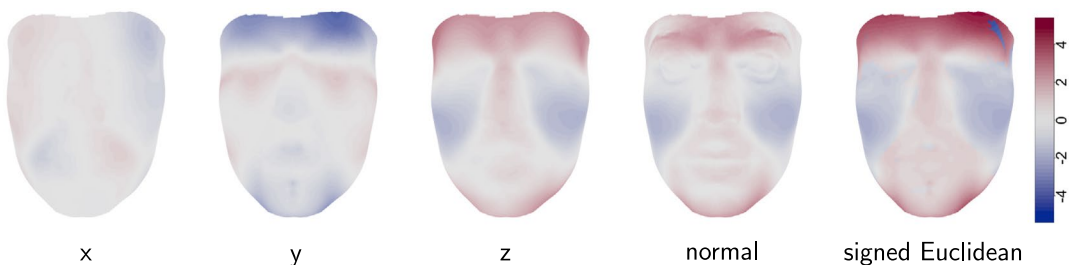

x   y   z   normal   signed Euclidean

FIGURE 4 The mean female face painted with colours to represent the movement (mm) required in different dimensions to reach the mean male face [Colour figure can be viewed at wileyonlinelibrary.com]

good place to start as it shows the broad, qualitative differences between the shapes, with the other options available as follow-up. For small-scale movements, the normal and signed-Euclidean distances are often effective and can give greater detail on the nature and size of the movement, as the two surfaces are generally close. Other plotting devices are available, such as the use of transparent surfaces, or one transparent surface with a wireframe representation. The best choice of display will depend on the particular features and differences of the shapes involved.

One of the most effective means of displaying differences is through animation, with the display of a sequence of intermediate steps along a path between the two shapes to be compared. Several of the plots in this paper have animated versions which are available in the *Supplementary information*.

3 | VISUALISING SHAPE DATA SETS

3.1 | Exploring variation

While a visual comparison of means is useful, an understanding of the variation involved in a data set is necessary for any form of statistical analysis. A simple device is to display the size of the variation at each location on the model. Figure 5 shows the value of $\log|\det(\widehat{\Sigma}_j)|$, where $\widehat{\Sigma}_j$ is the empirical covariance matrix of the x-, y- and z-coordinates at location j after Procrustes registration. The regions of higher variability include the eyes, whose reflective surface can introduce some inaccuracy, the forehead, which lies at the edge of the facial surface, and the chin and nasal tip, where the degree of prominence can vary considerably. Effects associated with the model curves which traverse the cheeks, where flatness can induce some variability in location, are also apparent.

Descriptions which capture the correlation between locations are clearly required. These also need to deal with the difficulty that the dimensionality of the shape representation (for example, 917 3D points in a discrete representation of a surface) often exceeds by a large margin the number of shapes present. Principal components offers an immediate route to the creation of a lower dimensional space which captures the principal features of shape variation. Procrustes registration places the aligned shapes X_i in a non-linear space but the *approximate tangent co-ordinates*, vec $(X_i - \overline{X})$, can be analysed very effectively as a linear space. Here the vec operator creates a vector of length $3J$ from
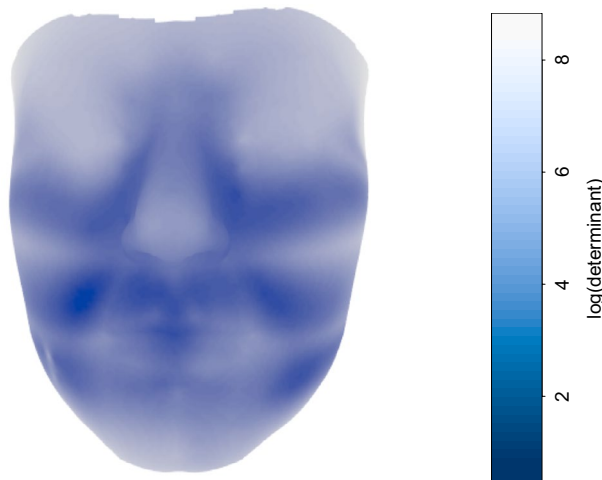

FIGURE 5 The variability in British female faces [Colour figure can be viewed at wileyonlinelibrary.com]

the $J \times 3$ matrix X by stacking its columns. The eigenvectors e_k and eigenvalues λ_k of the covariance matrix of the tangent co-ordinates then capture the directions along which variation is sequentially maximal. For the kth direction, a description of the variation involved is provided by considering $c\sqrt{\lambda_k}e_k$ which, when reassembled into a $J \times 3$ matrix using the vec^{-1} operator, represents shape variation from the mean along the principal component direction in multiples of the standard deviation. The multiplier c is often set at ± 2, or ± 3 if some magnification is required. Dryden and Mardia (2016, chapter 7) give all the details.

A functional data analysis perspective can be applied in this setting by following the pattern described by Ramsay and Silverman (1997). When the data are in the form of functional objects, $x(s)$, where s lies in an appropriate sample space $\mathcal{S}$, principal components are then defined as orthonormal functions $\beta(s)$ which successively maximise the variance of $\int \beta(s)x(s)ds$. In many settings, the sample space $\mathcal{S}$ is a time interval or a spatial region in standard Cartesian form. In the present setting, the functional object has the much more complex form of a 2D manifold embedded in 3D space. The immediate problem is how to parameterise this in a suitable sample space $\mathcal{S}$. A solution is provided by setting this to be the Procrustes mean shape, $\mathcal{S}_\mu$. Any other shape in the sample can then be expressed through the three functions $\{x(s), y(s), z(s)\}$ which give the 3D deviations of this shape from the mean at location s.

This takes us to the realm of multivariate functional principal components which seek to maximise the variance of

$$\int_{\mathcal{S}_\mu} \beta_x(s)x(s)ds + \int_{\mathcal{S}_\mu} \beta_y(s)y(s)ds + \int_{\mathcal{S}_\mu} \beta_z(s)z(s)ds,$$

as discussed by Ramsay and Silverman (1997, section 6.5.1). As usual, computations are conveniently based on discrete approximations to these integrals. The model form of each shape has a consistent triangulation so, for example, a convenient approximation can be written as

$$\int_{\mathcal{S}_\mu} \beta_x(s)x(s)ds \approx \sum_j \beta_x(s_j)x(s_j)a_j = \sum_j \left\{ \beta_x(s_j)\sqrt{a_j} \right\} \left\{ x(s_j)\sqrt{a_j} \right\},$$

where, as previously, j indexes the discrete surface representation and a_j is the area surrounding s_j. This allows functional principal components analysis to be implemented simply by applying standard principal components analysis to the discrete representations of a sample of n surfaces in the weighted form $\{(x_i(s_j)\sqrt{a_j}, y_i(s_j)\sqrt{a_j}, z_i(s_j)\sqrt{a_j}); i = 1, \ldots, n\}$. The square-root weights attached to the β's ensuring that the coefficient functions have norm (approximated integral of the square) equal to 1.

Figure 6 shows the results of applying functional principal components to the British female data. Given the high dimensionality of shape surfaces, the number of components required to capture a high proportion of the variation in the data may be reasonably large, with 9 components required to capture 82% of the variation in this case. Figure 6 shows the scores, $\text{vec}\,(X_i - \overline{X})^T e_k$, on these principal components, with the diminishing widths of the boxplots illustrating the gradual reduction in variation across the components. The shape changes associated with the first four principal components are indicated by superimposing the faces which correspond to $c = \pm 2$ standard deviations (pink and green). Methods for investigating individual components are discussed below in the context of comparing groups but the variation in a single group can be helpfully displayed through the idea of a 'grand tour', proposed for general multivariate data by Asimov (1985). A very simple version of this uses a vector of p independent normal random variables r to create a random sample of locations in the space of the first p principal components, $\{r_k\sqrt{\lambda_k}e_k; k = 1, \ldots, p\}$. Turning these into shapes and tracking between
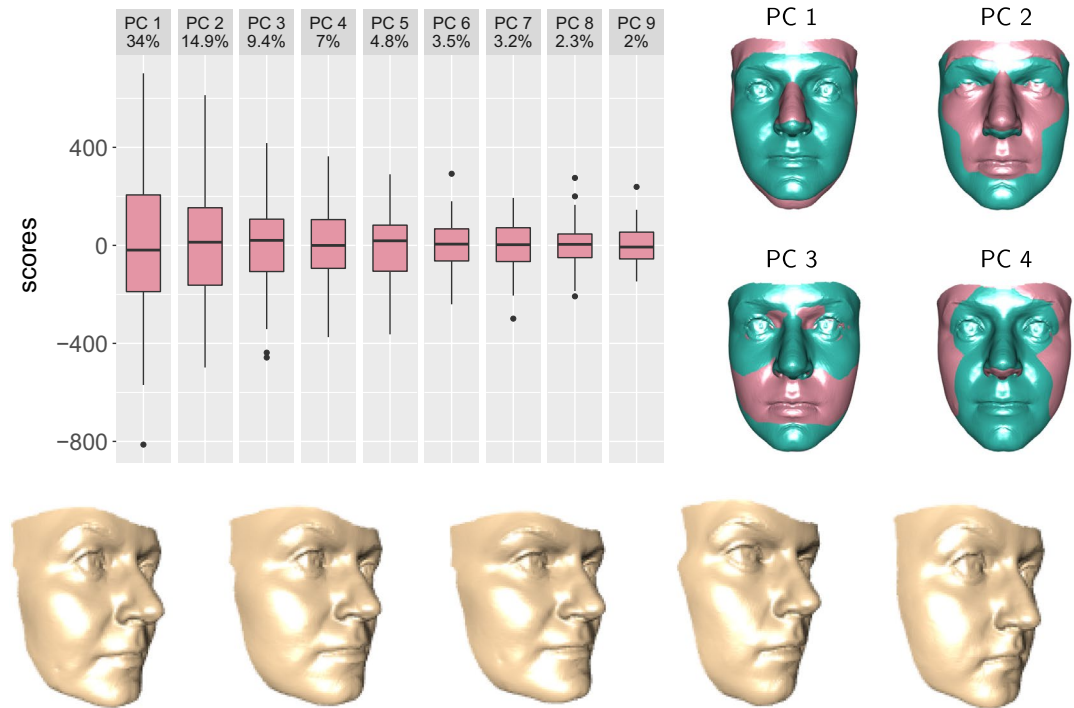
FIGURE 6 The top left-hand plot shows the scores for the first 10 principal components for British females. The four top right-hand plots show the nature of the shape change associated with the first 4 principal components. The lower plots show randomly generated faces from a 'grand tour' of the variation in British female faces. See the *Supplementary information* for animations [Colour figure can be viewed at wileyonlinelibrary.com]

successive pairs by simple interpolation creates an animation which randomly explores the variation in shape. Figure 6 illustrates five random positions which form the staging posts of a tour and so provides a visualisation of the range of shapes exhibited in the population of interest. This approach forms the basis of a comparison between individuals and a control data set in Section 4.

3.2 | Assessing differences between groups

When groups representing different populations are present in a data set, Procrustes registration is usually applied to the whole data set, to create a common shape space. (This is in contrast to the simple illustration used in Figure 3.) Principal components provide a helpful way of reducing the dimensionality of the space in which comparison takes place, while retaining as much of the variability as possible. If components are constructed from the combined data set, without reference to the group structure, then the variation captured by each component will contain both intra- and inter-group contributions. The top left-hand plot in Figure 7 displays the scores for the principal components constructed in this way for the sexual dimorphism data. Over 80% of the variation is captured by 10 principal components. As the signs of the eigenvectors which define the components are arbitrary, these have been reversed where necessary to ensure that the male mean score is higher, for ease of interpretation. It is often the case that the first few components capture large scale variation (greater width and smaller height etc.) which is common across groups, with group differences associated with more subtle aspects of shape.
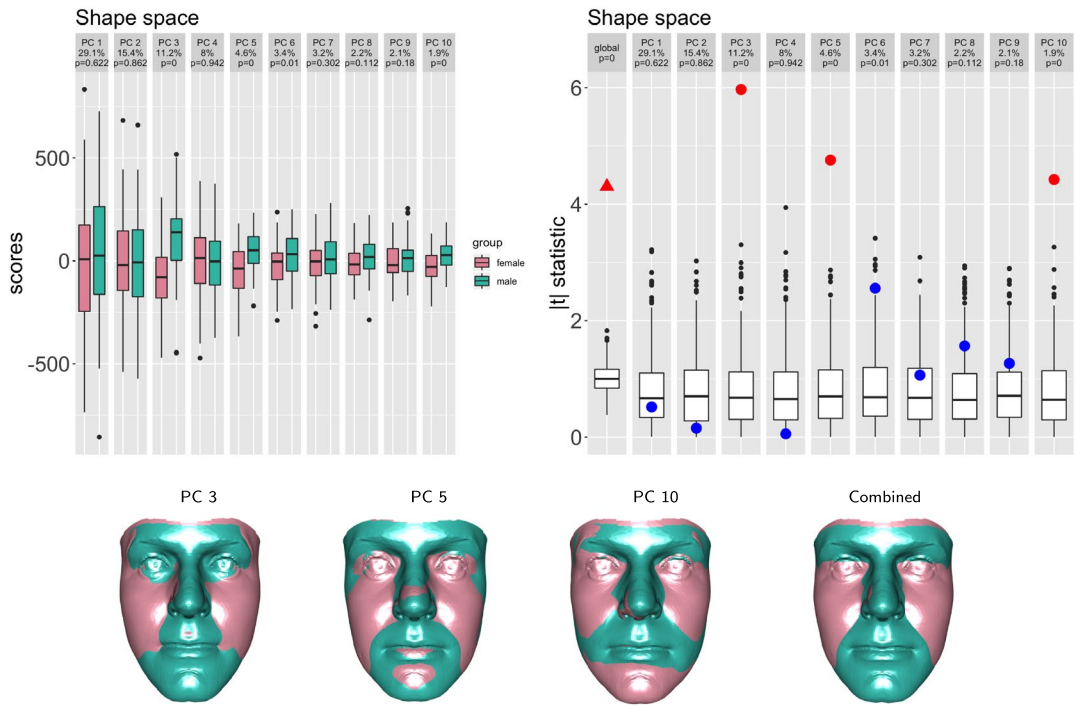
FIGURE 7 The top row shows the scores for males and females for the first 10 principal components constructed in tangent space and the simulated test statistics (boxplots) and observed values (triangle and dots) for a test of identical distributions. Test statistics showing evidence for differences are plotted in red. The second row shows the shape change associated with each significant component and their combination. Green and pink refer to the male and female ends of the scale respectively [Colour figure can be viewed at wileyonlinelibrary.com]

In the reduced space of the first p components, a global assessment of the evidence for mean differences in male and female shape is provided by Hotelling's T^2 statistic, $T^2 = \frac{1}{(1/n_m + 1/n_f)} \left(\bar{v}_m - \bar{v}_f\right)^T \hat{\Sigma}^{-1} \left(\bar{v}_m - \bar{v}_f\right)$, where $\bar{v}_m$, $\bar{v}_f$ denote the mean p-dimensional score vectors, n_m, n_f denote the group sample sizes and $\hat{\Sigma}$ denotes the usual estimate of the common covariance matrix of the scores. It is also tempting to explore the nature of any evidence of differences by examining the t-statistics, $\frac{(\bar{v}_{km} - \bar{v}_{kf})}{\sqrt{(1/n_m + 1/n_f)}\hat{\sigma}_k}$, where $\hat{\sigma}_k$ denotes the estimate of the common standard deviation of the groups on the kth component. However, there are concerns about the appropriateness of the usual reference distributions for these statistics, particularly in a situation where dimensionality exceeds sample size. Bedrick (2019) also shows that even in the 'non-sparse' case the distributional properties of test statistics can be affected by the construction of the component directions in terms of optimising variance, under some circumstances. In light of both of these issues, a permutation approach is attractive. Here the reference distributions for the test statistics are generated empirically simply by recomputing the values from the data set with group labels randomly permuted 500 times. The top right-hand panel of Figure 7 illustrates this. In order to make the distribution for the global test comparable with the others, the values of $\sqrt{T^2/p}$ have been plotted and the individual t-statistics have been placed on an absolute value scale. The values of the test statistics computed from the observed data are shown as a triangle (global) and dots (components), with the empirical p-values noted at the top of the plot. The global test gives strong evidence of differences in means while components 3, 5 and 10 are identified as strong sources of difference. To adjust for the multiple comparisons involved, a Bonferroni threshold for the empirical p-values has been adopted as $0.05/10 = 0.005$ and colour (red) has been used to indicate where that threshold has

been exceeded. As Bedrick (2019) points out, it is important to note that the *p*-values associated with individual components should be interpreted in the context of the null hypothesis that the mean scores are identical for *all* components simultaneously.

The images in the lower part of Figure 7 indicate, for those components which exhibit evidence of differences in groups, the nature of the associated shape change in the usual form of $\pm 2\sqrt{\lambda_k}$ from the mean. The image corresponding to the positive end of the scale is more strongly associated with male shape (green) and the negative end of the scale with female shape (pink). The association of male shape with more prominent nose, chin and eyebrows, and female shape with more prominent cheeks, is apparent. However, the individual components cannot be given special status in the description of male-female differences as they were constructed simply by maximising the variation explained across the whole data set. The strongest differences may in fact lie in directions which are combinations of components. In order to explore the nature of shape differences, following a significant global T^2 test, it is therefore helpful to construct a combined display which corresponds to movement simultaneously along all the components which exhibit evidence of differences. This is aided by the earlier modification of components to ensure that positive signs are more strongly associated with males. Movement to $\pm 2\sqrt{\lambda_k}$ in all components simultaneously would construct a rather extreme shape so the values $\pm 2\sqrt{\lambda_k}/\sqrt{q}$ are used, where q denotes the number of components in simultaneous movement (here $q = 3$). This ensures that the resulting shape sits on the same quantile contour of a multivariate normal distribution as the shapes which move the individual components to $\pm 2\sqrt{\lambda_k}$. The result for the sexual dimorphism data is shown at the right-hand end of the row of facial images in Figure 7. This gives a very helpful representation of the combined effects of the individual components which carry evidence of differences. It also has the attractive property of giving stronger weight to those components which explain larger amounts of variation. The overall difference in shape change is clear and corresponds closely to the comparison of means in Figure 3 but this is now backed up by convincing statistical evidence.

When group differences are of interest, an alternative approach to principal components is through the intra-group covariance. As pointed out by Dryden and Mardia (2016, chapter 9), the T^2 statistic can be written as

$$T^2 = \sum_{k=1}^{p} \left[\frac{\bar{v}_{k1} - \bar{v}_{k2}}{\sqrt{\lambda_k (1/n_1 + 1/n_2)}} \right]^2,$$

where the λ_k denote the eigenvalues and $\bar{v}_{k1}$, $\bar{v}_{k2}$ the mean principal component scores, using the eigenvectors derived from an estimate of the common covariance matrix. Dimensionality reduction follows from the truncation to p terms, with each individual term having the attractive interpretation of the square of a two-sample *t*-statistic on the scores from each component. The warnings about distributional properties again apply, with a permutation test providing a convenient solution. However, this time the eigen decomposition needs to be performed for every random permutation because estimation of the common covariance matrix depends on the group structure.

Figure 8 shows the results of this 'group shape space' approach on the sexual dimorphism data, with $p = 10$ to remain consistent with the earlier example. The absolute value scale has been used again for the boxplots. This loses the property that the global statistic is a simple average of its components but the *t*-statistic scale is helpful, and there is no effect on the performance of the tests. The smaller facial images show the nature of the shape change associated with the individual components (3, 5, 7) where there is strong evidence of differences between males and females. There is no reason why the differences in mean shape should align with the axes of the common covariance matrix so,
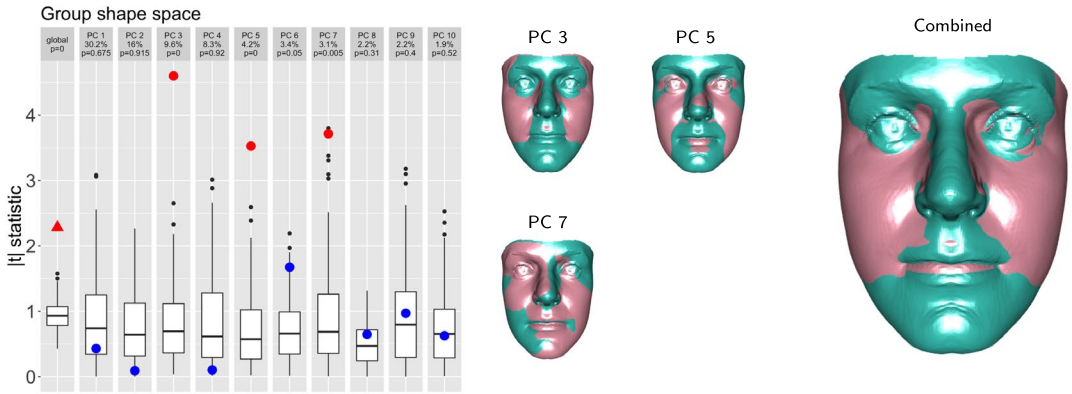
FIGURE 8 The left-hand plot shows the simulated test statistics (boxplots) and observed values (triangle and dots) for a test of identical male and female distributions, using the components derived from the common covariance matrix. The small faces shows the shape change associated with each significant component and the larger face illustrates their combined effect [Colour figure can be viewed at wileyonlinelibrary.com]

again, the individual components do not have special status. The larger facial image shows the shape change associated with the combination of these three components. This characterises the sub-space where the evidence for difference is strongest and it is reassuring to see that this is very similar to the sub-space identified from the principal components which do not exploit group structure, as displayed in Figure 7. This underlines the case for identifying and interpreting the sub-space as a whole, with the components simply providing particular indexing bases.

3.3 | Affine/non-affine decomposition

A further sub-space approach is available through partitioning the variation in the data into affine and non-affine components. The former involves linear transformations which apply across the whole object of interest. The latter space contains non-linear transformations which describe local and more complex effects. Rohlf and Bookstein (2003) showed that these sub-spaces can be easily created from the Procrustes aligned shapes $\{X_i; i = 1, \ldots, n\}$ through the regression models

$$X_i = \overline{X}\alpha_i + \varepsilon_i,$$

where the α_i denote 3×3 matrices of regression coefficients. The affine co-ordinates X_{Ai} are then available as the fitted values while the non-affine co-ordinates X_{Ni} are obtained by adding the residuals to the mean as

$$X_{Ai} = \overline{X}\widehat{\alpha}_i,$$
$$X_{Ni} = \overline{X} + (X_i - X_{Ai}),$$

where $\widehat{\alpha}_i$ denotes the least squares estimates. More formally, the algebra associated with linear regression, particularly the independence of residuals and fitted values, separates the space of the Procrustes shape co-ordinates X_i into two orthogonal sub-spaces which capture the affine and non-affine behaviours. Analysis can therefore proceed separately within these sub-spaces to provide complementary descriptions of the variation in the data set.

The comparison of British and Chinese female facial shapes provides a simple example. Visual discrimination between these two ethnic groups is usually straightforward but examination of mean shapes allows the distinctive features to be identified and quantified. Figure 9 shows the nature of shape change in different sub-spaces, based on a sample of 69 British and 26 Chinese adult subjects. In overall tangent space, there is only one significant component, with the associated shape change displayed in the central image in the usual form of ± 2 standard deviations around the mean. This draws attention to the more prominent central areas in British subjects (pink), including the brow ridge, nose, oral region and chin. Correspondingly, the overall shape of Chinese faces (green) is flatter than that of their British counterparts, with more prominent cheek areas. This is an example where intergroup differences dominate the variation in the data. Indeed, the strength of this difference is indicated by the fact that there is no overlap between the scores of the British and Chinese faces on the first principal component in tangent space (not shown).

It is interesting to explore whether these differences can be explained by affine transformation or whether non-affine transformations are required. In the affine sub-space, only the first principal component shows clear evidence of difference between the groups and it is already clear that the lower brow of the British faces is not captured in this sub-space. This is confirmed by analysis in the non-affine sub-space where there are two principal components which exhibit clear evidence of differences between the groups and whose combined effects are displayed in the right-hand image of Figure 9.

4 | VISUALISING THE SHAPE OF INDIVIDUALS

In the previous section, evidence for systematic differences between groups was considered, while allowing for the presence of individual variability. This section considers situations where interest lies in the evaluation of individuals. Traits which can be expressed in single values are considered, as well as more general characterisation of the particularities of individual shapes.

4.1 | Asymmetry

For shapes whose ideal form is symmetric, deviations which disturb this symmetry are important features. The left/right symmetry of human faces is a major example, where any strong departure from symmetry creates a striking visual impression. However, real faces are all asymmetric to some degree so, as part of the process of evaluating an individual shape, it is important to characterise the asymmetries found in an appropriate reference population.

When a shape is represented by a set of point locations, some of which are paired as left/right counterparts, quantification of asymmetry is generally based on the degree of post-registration mismatch
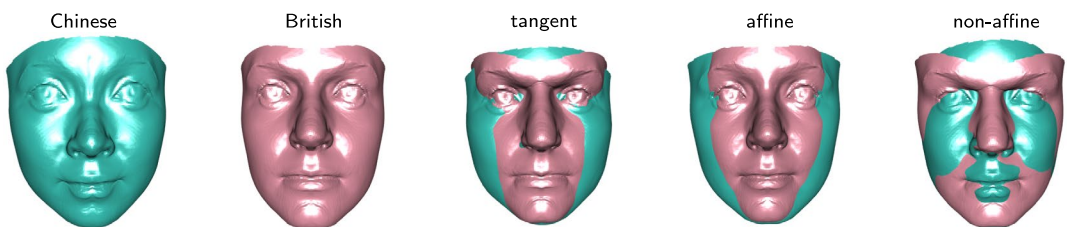

FIGURE 9   From left to right, the images show the mean Chinese female face, the mean British female face, and the combined principal components of shape change in tangent, affine and non-affine spaces respectively [Colour figure can be viewed at wileyonlinelibrary.com]

between the shape and its reflection with the left/right labels swapped. Theoretical development of this idea was undertaken by Mardia, Bookstein and Moreton (2000) and Kent and Mardia (2001) in the context of landmarks, and many authors have exploited this thinking in biological contexts. Bock and Bowman (2006) proposed a decomposition of global asymmetry which allowed local sources to be identified and separated into contributions from individual features and their configurations.

The first step in computing a functional measure of asymmetry for a surface $X(s)$ is to apply (functional) Procrustes matching of the mirror image onto the original surface, to create the new surface $\tilde{X}(s)$. The mirror image is created in practice by reflecting and relabelling the configuration of points which express the shape model. The integrated comparison and its discrete approximation are then easily constructed as

$$\frac{1}{A(\mathcal{S})}\int_{\mathcal{S}}||X(s)-\tilde{X}(s)||^2 \mathrm{d}s \approx \frac{1}{\left(\sum_{j=1}^{J} a_j\right)}\sum_{j=1}^{J}||X(s_j)-\tilde{X}(s_j)||^2 a_j,$$

where, to be even-handed, $\mathcal{S}$ is the surface formed from the average of $X(s)$ and $\tilde{X}(s)$, $A(\mathcal{S})$ denotes its surface area and, as usual, the a_j give the areas surrounding the discrete surface points s_j. The final asymmetry score is achieved by applying a square root transformation, so that the scale of the end result matches the scale of the original co-ordinate measurements.

Earlier applications in the paper have been based on human sexual dimorphism but the new context of orthognathic surgery in now introduced. This involves the repositioning of the maxilla or mandible of a patient to improve the alignment of teeth and to address issues of facial appearance. Vittert et al. (2018) included asymmetry in their assessment of post-surgical outcome, identifying evidence of a reduction in mean asymmetry in the upper lip region. However, asymmetry scores also give the opportunity to assess patients individually. Figure 10 shows the facial image of one post-surgical patient together with comparison of the reflected and matched image in both superimposition and colour-coded forms. Interpretation of this information is informed by quantifying the asymmetry scores exhibited in an adult control population, with the distributions represented in the right-hand side of Figure 10 by density strips (Jackson, 2008). The global asymmetry scores, for both pre-surgical and post-surgical facial shapes of this patient, have been superimposed on the bottom density strip. These scores are entirely typical of controls and in particular they provide reassurance that surgery has not introduced any marked asymmetry overall. The scores have also been computed for a variety of sub-regions, indicated in the top right-hand image of the four facial images in Figure 10. The scores and density strips indicate strong nasal asymmetry, but this is apparent both before and after surgery and so it cannot be attributed to surgical intervention.

4.2 | Closest controls

For more general assessment of the shapes of individual cases, an approach analogous to the concept of a 'normal range' for univariate data is required. In the surgical context, characterising any differences between a post-surgical patient and a control population could provide helpful guidance on the nature of any further surgery which may be required. Bowman and Bock (2006) outlined an approach based on the concept of a 'closest control'. This identifies the shape which is as close as possible to the individual of interest but which lies on the surface of a 95% prediction ellipsoid and so lies within the 'normal range' associated with controls. Any remaining shape differences then characterise the features of the individual shape which are different from controls.
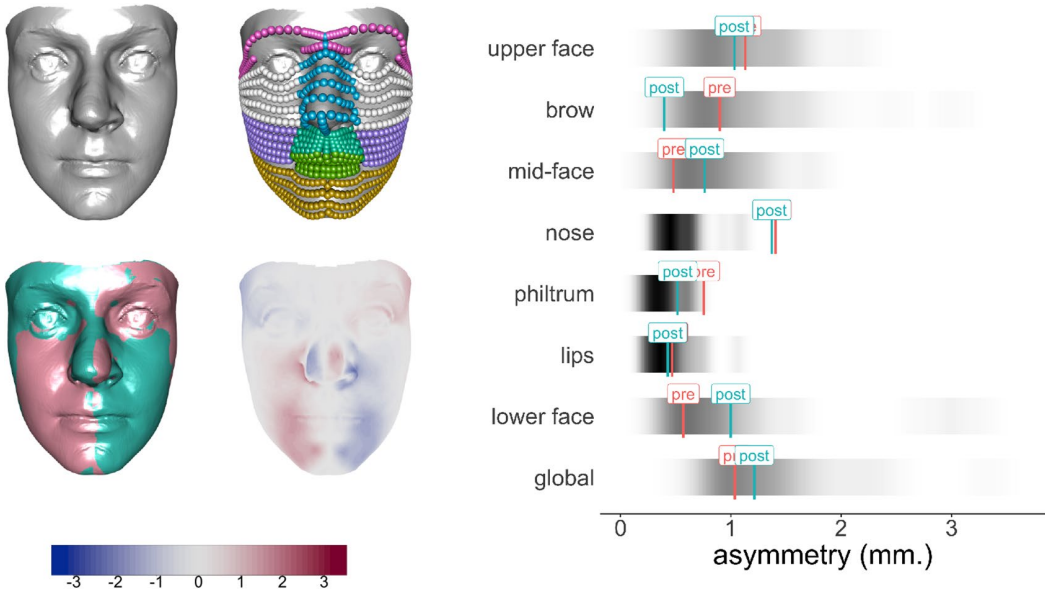
FIGURE 10   The four facial images show, in clockwise order from top left, a post-surgical patient, the sub-regions used to compute asymmetry scores, the colour-coded distance between the shape and its matched reflection, and the superimposition of the shape and its matched reflection. The density strips show the asymmetry scores from control faces, with the pre- and post-surgical scores for the individual superimposed. (The facial images have been anonymised by warping to a reference shape.) [Colour figure can be viewed at wileyonlinelibrary.com]

Bowman and Bock (2006) derived the algebra of this in a simple case involving curve data, using a principal component regularisation to reduce dimensionality across both cases and controls. The concept is applied here to surface data but the ideas are developed further in two important ways. Firstly, principal components are constructed from the control data only. This gives a clear interpretation of the components which is unaffected by the particular cases available. Secondly, variation unexplained by these components is also considered, in order to give a complete description of the observed data.

If a new shape Z, such as a post-surgical patient, is registered onto the control mean then it can be projected into the space of the first p principal components, denoted by C_p, by computing the score vector $v(Z) = \text{vec} \ (Z - \overline{X})^T E_p$, where E_p is the matrix whose columns contain the first p principal component vectors derived from the control data. Within this space, the Mahalanobis distance of the new shape from the mean control shape is

$$d(Z) = v(Z)^T \widehat{\Sigma}^{-1} v(Z)$$

where $\widehat{\Sigma}$ is a diagonal matrix containing the variances of the principal components. The Mahalanobis distance has a χ_p^2 distribution approximately. If $d(Z)$ is less than the 95th percentile of this distribution, denoted by $\chi_p^2(0.95)$, then the new shape falls within the 'normal range' of controls in this space. If $d(Z) > \chi_p^2(0.95)$ then the *closest control* in this p-dimensional space can be found by shrinking v towards 0 until its Mahalanobis distance matches $\chi_p^2(0.95)$. The shrinking factor α_1 is easily found as $\alpha_1 = \sqrt{\frac{\chi_p^2(0.95)}{d(Z)}}$, by solving the equation $\alpha_1 v(Z)^T \widehat{\Sigma}^{-1} \alpha_1 v(Z) = \chi_p^2(0.95)$. The scores of this new location $\alpha_1 v(Z)$ are then converted into tangent co-ordinates as $\alpha_1 v(Z) E_p^T$, and expressed as a shape by reconfiguring the tangent co-ordinates into a three column matrix in the usual manner as

$$cc_p(Z) = \overline{X} + \ \text{vec}^{-1} \{\alpha_1 v(Z) E_p^T\}.$$

This finds the closest control in C_p. However, the case of interest may well have shape features which cannot be captured in this space so characterisation in the complementary space, denoted by $\mathcal{R}_p$, is also required. The projection of Z onto C_p is

$$\tilde{Z} = \text{vec}^{-1}\left\{ v(Z)E_p^T \right\},$$

so the relevant information is found in the residual shape $R(Z) = Z - \tilde{Z}$. The length of the residual at each model location can be quantified in the vector $L(Z) = \sqrt{R(Z)^2 1_3}$, where here the square-root and the exponent 2 are applied element-wise. A measure of variation in the lengths of the residuals at each model location for controls is then available in the vector v whose jth element is the standard deviation of $\{L(X_i)_j; i = 1, \ldots, n\}$, where $L(X_i)_j$ denotes the jth element of $L(X_i)$. A simple measure of variation is then

$$r(Z) = \frac{1}{J} \sum_{j=1}^{J} L(Z)_j / v_j.$$

This averages the lengths of the residuals across the model locations, standardised at each model location by the variation in control residual length. The value of $r(Z)$ may be regarded as atypical if it lies beyond q_{95}, the 95th percentile of $\{v(X_i); i = 1, \ldots, n\}$. A closest control in the residual space, R_p, can then be constructed by shrinking the residual shape to $\alpha_2 Z$, where $\alpha_2 = q_{95}/r(Z)$. An overall closest control for z can now be constructed as

$$cc(Z) = cc_p(Z) + \alpha_2 R(Z)$$

which combines the closest controls in the sub-spaces R_p and C_p.

Figure 11 shows the results of applying the concept of closest control to two post-surgical orthognathic cases. The left-hand histogram shows the Mahalanobis distances of controls in C_p, the space of the first p principal components for controls. The use of $p = 9$ was determined by the smallest number of components which explained at least 80% of the variation in the controls. Case 1 clearly lies in the tails of the control distribution while case 2 is unexceptional. The right-hand histogram shows that both cases exhibit unusual behaviour in the residual space, $\mathcal{R}_p$. However, shape differences in this residual space may be small. The facial images in the lower part of Figure 11 compare case 1 (green) with its closest control (pink) by superimposition and by normal distances. (The facial images have been anonymised by warping to a reference shape while preserving the contrast with the closest control.) This characterises the unusual features of the case as a slightly more prominent lower face than in controls, particularly in the mandible (lower jaw). This is potentially valuable feedback on surgery which involves repositioning of the underlying bones. A display of the closest control information for case 2 is deferred to the next sub-section.

4.3 | An integrated patient assessment

The methods described in this section provide valuable tools for the characterisation of individual shapes of interest. The combination of these tools forms the basis of an integrated patient assessment. This is illustrated in Figure 12, using case 2 from Figure 11. This brings together the observed pre-surgical and post-surgical shapes, comparisons of this case with control shape both for closest control analysis and for asymmetry, and illustrates differences in shape through superimposition and normal distances. (The facial images have again been anonymised by warping to a reference shape.)
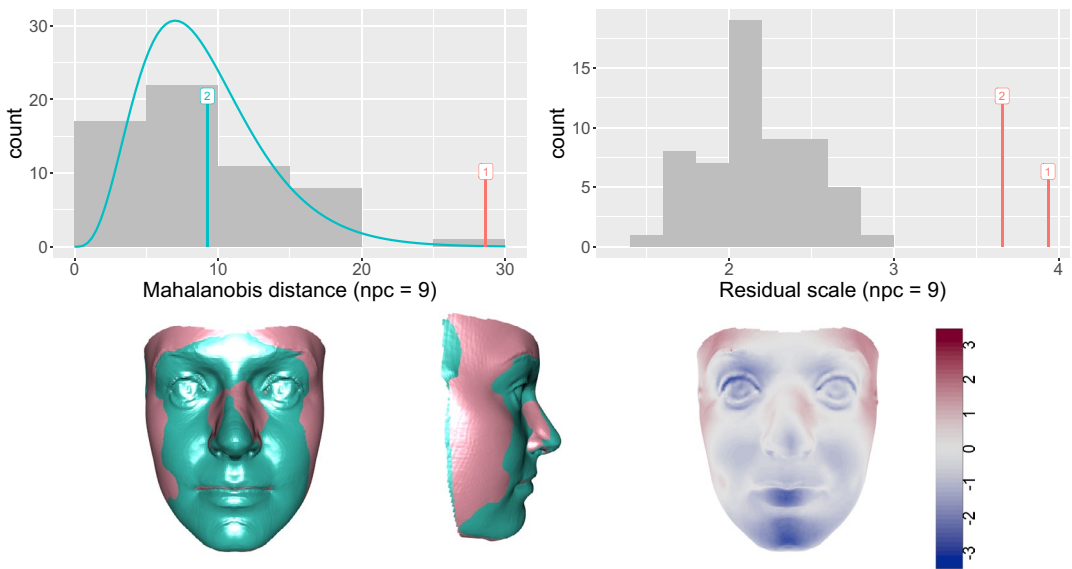
FIGURE 11 The histograms show the distances of the control shapes from the mean in the space of the first 9 principal components (left) and in the residual space (right). The distances of two post-surgical cases from the control mean are superimposed. The lower images compare the facial shape of case number 1 (green) with its closest control, both through superimposition and as normal distances from case to closest control. (The facial images have been anonymised by warping to a reference shape.) [Colour figure can be viewed at wileyonlinelibrary.com]

An interactive display would allow those reviewing the case to inspect the shapes in 3D and to query further information. However, this static display gives a helpful summary of the effects of surgery on this particular patient.

5 | DISCUSSION

This paper has proposed methods of analysis for high resolution surface data and corresponding models which give consistent descriptions of each observed shape. A strong emphasis has been on the adoption of functional forms of analysis and practical implementation has been based on the application of standard forms of analysis to suitably weighted versions of the discrete representations of each surface. Given the high dimensionality of these surface representations, more sophisticated methods of analysis such as those based on spiked covariance models (Johnstone, 2001) could also be applied.

Some forms of analysis required identification of a common sample space and this was overcome by using the mean surface as an indexing shape. This enabled functional forms of registration, principal components analysis and group comparisons to be developed. In applying these methods, strong emphasis was also placed on the use of principal components to identify sub-spaces of interest rather than inspection of individual components. In some applications, there may be prior information on particular regions where systematic differences are more likely to occur and this could be used to guide the particular sub-spaces examined. At the very simplest level, the methods described in the paper could be applied to particular sub-regions of the overall surface.

Tests have been applied to assess the evidence for shape differences and support an exploration of the nature of these differences. The use of a permutation approach should ensure that the size of
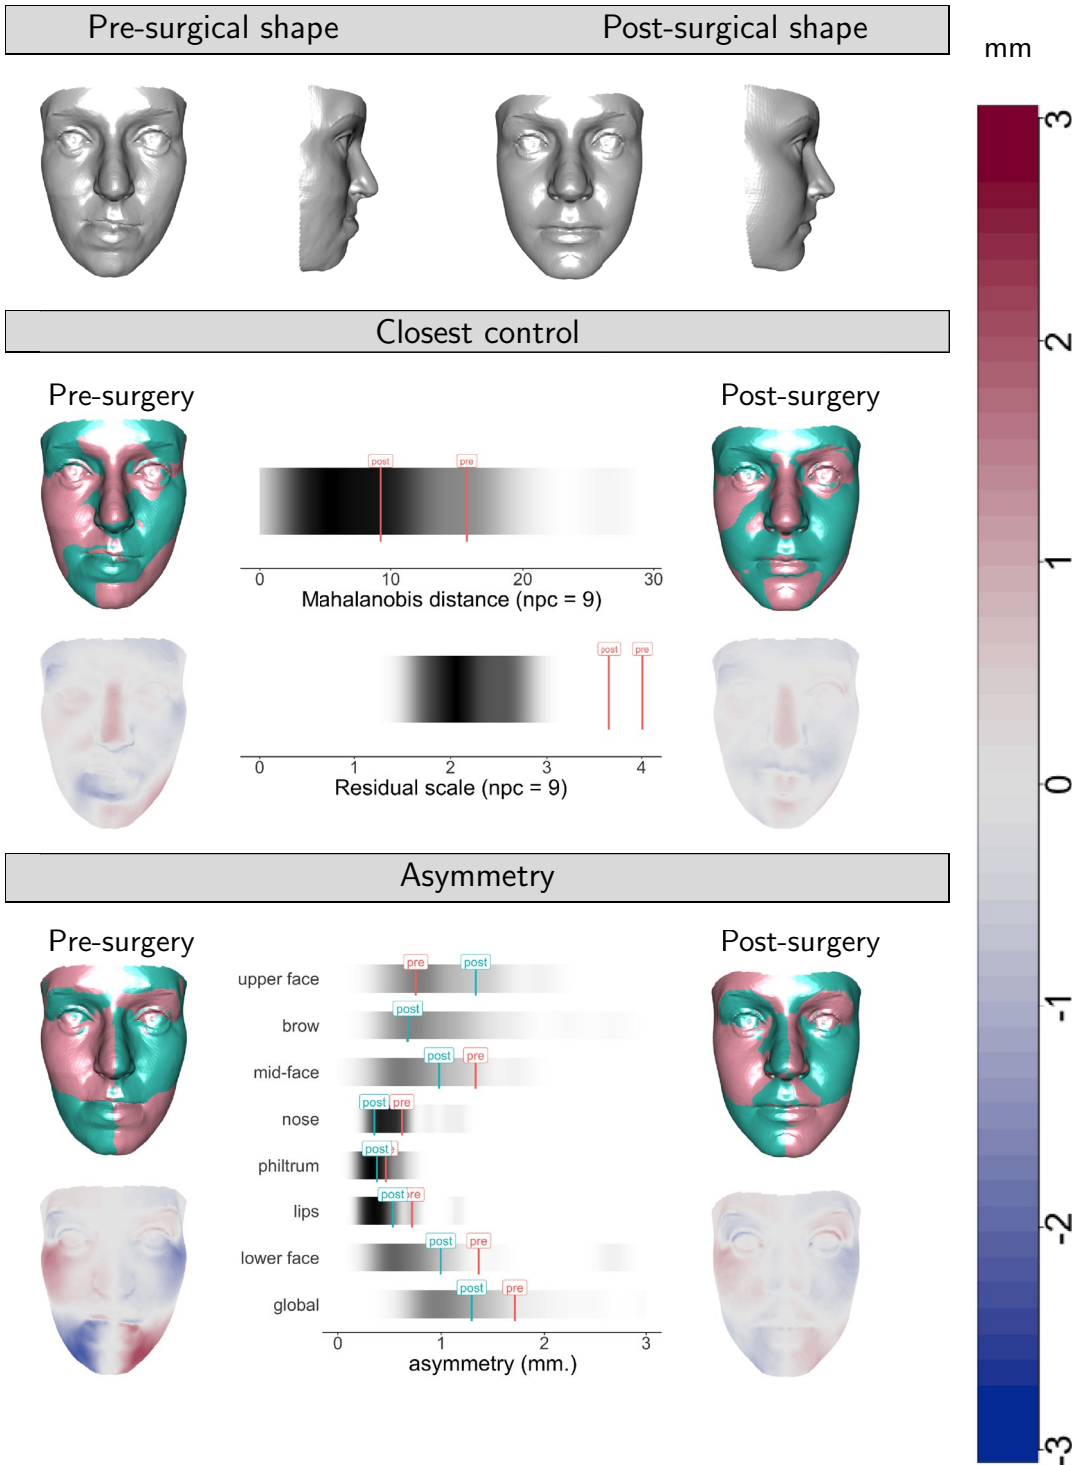
FIGURE 12 An integrated assessment of patient 2 from Figure 11. The facial images have been anonymised by warping to a reference shape [Colour figure can be viewed at wileyonlinelibrary.com]

the tests are correct but an understanding of power would be helpful. Simulation studies would be an obvious way of investigating this.

Graphical displays of shape differences were also described. Particular attention was given to the comparison of individual shapes with a relevant control group. In additional to univariate measures such as asymmetry scores, the concept of a 'closest control' was developed in detail, to give a means of identifying any unusual characteristics of an individual shape of interest. We believe this to be a potentially very powerful approach and further practical implications of its use, particularly in surgical contexts, will be the subject of subsequent research.

REFERENCES

Adler, D. & Murdoch, D. (2019) rgl: 3D Visualization Using OpenGL.

Armann, R. & Balthoff, I. (2012) Male and female faces are only perceived categorically when linked to familiar identities and when in doubt, he is a male. *Vision Research*, 63, 69–80.

Asimov, D. (1985) The grand tour: A tool for viewing multidimensional data. *SIAM Journal on Scientific and Statistical Computing*, 6(1), 128–143.

Bedrick, E.J. (2019) Data reduction prior to inference: Are there consequences of comparing groups using a t-test based on principal component scores? *Biometrics,*, 76, 508–517.

Bock, M.T. & Bowman, A.W. (2006) On the measurement and analysis of asymmetry with applications to facial modelling. *Applied Statistics*, 55, 77–91.

Bookstein, F.L. (1989) Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6), 567–585.

Bookstein, F.L. (1997) *Morphometric tools for landmark data: Geometry and biology*. Cambridge: Cambridge University Press.

Bowman, A.W. & Bock, M.T. (2006) Exploring variation in three-dimensional shape data. *Journal of Computational and Graphical Statistics*, 15(3), 524–541.

Brignell, C., Dryden, I. & Browne, W. (2015) Covariance weighted procrustes analysis. *Riemannian computing in computer vision*. Berlin: Springer, pp. 189–209.

Bruce, V., Burton, A.M., Hanna, E., Healey, P., Mason, O., Coombes, A.. et al. (1993) Sex discrimination: How do we tell the difference between male and female faces? *Perception*, 22(2), 131–152.

Claes, P., Walters, M., Shriver, M.D., Puts, D., Gibson, G., Clement, J. et al. (2012) Sexual dimorphism in multiple aspects of 3d facial symmetry and asymmetry defined by spatially dense geometric morphometrics. *Journal of Anatomy*, 221(2), 97–114.

Dryden, I.L. & Mardia, K. (2016) *Statistical shape analysis, with applications in R*, 2nd edition. New York: Wiley.

Duchon, J. (1977) *Constructive theory of functions of several variables*. Berlin: Springer, pp. 85–100.

Frelat, M.A., Katina, S., Weber, G.W. & Bookstein, F.L. (2012) A novel geometric morphometric approach to the study of long bone shape variation. *American Journal of Physical Anthropology*, 149(4), 628–638.

Gunz, P., Mitteroecker, P. & Bookstein, F.L. (2005) Semilandmarks in three dimensions. *Modern morphometrics in physical anthropology*. Berlin: Springer, pp. 73–98.

Hammond, P., Hutton, T.J., Allanson, J.E., Campbell, L.E., Hennekam, R., Holden, S. et al. (2004) 3d analysis of facial morphology. *American Journal of Medical Genetics Part A*, 126(4), 339–348.

Jackson, C.H. (2008) Displaying uncertainty with shading. *The American Statistician*, 62(4), 340–347.

Johnstone, I.M. (2001) On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29(2), 295–327.

Kent, J.T. & Mardia, K.V. (2001) Shape, procrustes tangent projections and bilateral symmetry. *Biometrika*, 88(2), 469–485.

Koenderink, J. (1990) *Solid shape*. 2, Cambridge: Cambridge Univ Press.

Koenderink, J. & van Doorn, A. (1992) Surface shape and curvature scales. *Image and Vision Computing*, 10(8), 557–564.

Mao, Z., Ju, X., Siebert, J.P., Cockshott, W.P. & Ayoub, A. (2006) Constructing dense correspondences for the analysis of 3d facial morphology. *Pattern Recognition Letters*, 27(6), 597–608.

Mardia, K.V., Bookstein, F.L. & Moreton, I.J. (2000) Statistical assessment of bilateral symmetry of shapes. *Biometrika*, 87(2), 285–300.

Mitteroecker, P. & Bookstein, F. (2008) The evolutionary role of modularity and integration in the hominoid cranium. *Evolution: International Journal of Organic Evolution*, 62(4), 943–958.

Mitteroecker, P., Gunz, P., Bernhard, M., Schaefer, K. & Bookstein, F.L. (2004) Comparison of cranial ontogenetic trajectories among great apes and humans. *Journal of Human Evolution*, 46(6), 679–698.

Paulsen, R.R. & Hilger, K.B. (2003) Shape modelling using markov random field restoration of point correspondences. *Biennial International Conference on Information Processing in Medical Imaging*. Springer, pp. 1–12.

Ramsay, J.O. & Silverman, B.W. (1997) *Functional Data Analysis*. New York: Springer.

R Core Team. (2019) *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Reeger, J., Fornberg, B. & Watts, M. (2016) Numerical quadrature over smooth, closed surfaces. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2194), 20160401.

Rohlf, F.J. & Bookstein, F.L. (2003) Computing the uniform component of shape variation. *Systematic Biology*, 52(1), 66–69.

Thompson, D.W. (1917) *On Growth and Form*. Cambridge, UK: Cambridge University Press.

Vittert, L., Bowman, A.W. & Katina, S. (2019) Statistical models for manifold data with applications to the human face. *Annals of Applied Statistics*, 13(4), 2539–2563.

Vittert, L., Katina, S., Ayoub, A., Khambay, B. & Bowman, A. (2018) Assessing the outcome of orthognathic surgery by three-dimensional soft tissue analysis. *International Journal of Oral and Maxillofacial Surgery*, 47(12), 1587–1595.

Waddington, J.L., Katina, S.K., O'Tuathaigh, C.M.P. & Bowman, A.W. (2017) Translational genetic modelling of 3d craniofacial dysmorphology: Elaborating the facial phenotype of neurodev elopmental disorders through the "prism" of schizophrenia. *Current Behavioral Neuroscience Reports*, 4(4), 322–330.

Wilkinson, C. (2004) *Forensic facial reconstruction*. Cambridge: Cambridge University Press.

Zeileis, A., Hornik, K. & Murrell, P. (2009) Escaping rgbland: Selecting colors for statistical graphics. *Computational Statistics & Data Analysis*, 53(9), 3259–3270.

SUPPORTING INFORMATION

Animated versions of some of the figures in the paper are available by downloading files from http://dx.doi.org/10.5525/gla.researchdata.1130.

APPENDIX A

SURFACE CURVATURE

The local shape at a 3D location $m = (x, y, z)$ on a differentiable manifold M can be characterised through the quadratic surface

$$z = \frac{1}{2} \left(\kappa_1(m)x^2 + \kappa_2(m)y^2 \right),$$

where z lies in the normal direction to the surface at *m* and the orthogonal axes x and y lie on the tangent plane, associated with the directions of maximum, $\kappa_1(m)$, and minimum, $\kappa_2(m)$, curvature. This is clearly described by Koenderink (1990) and many others, along with a wide variety of other key tools for studying surface shape. The coefficients, $\kappa_1(m)$ and $\kappa_2(m)$, are referred to as the *principal curvatures* and, along with their associated *principal directions*, d_1 and d_2, they provide the essential information for characterising curvature across the manifold. A wealth of further detail is available in Koenderink (1990).

There is a variety of summary measures of curvature. An important one is *Gaussian curvature*, defined as the product of the principal curvatures, $\kappa_1(m)\kappa_2(m)$, which provides a summary of the size of curvature present at each location. Another particular quantity of interest is the *shape index*, which characterises the type of curvature present at each location. This is defined as

$$S(m) = \frac{2}{\pi}\tan^{-1}\left(\frac{\kappa_2(m) + \kappa_1(m)}{\kappa_2(m) - \kappa_1(m)}\right).$$

Koenderink and van Doorn (1992) provide further details.

APPENDIX B

3D WARPING

The technical details of warping are described here because the method is not widely used in 3D. We seek a function which maps X onto Y exactly. If an interpolant of a single co-ordinate of Y as a function of the three co-ordinates of X is considered, then the elegant functional analysis described by Duchon (1977) provides an immediate solution. The aim is to find the interpolating function f which has minimal bending energy, defined as

$$\int_{\mathcal{R}^3}\left\{\sum_{p=1}^{3}\sum_{q=1}^{3}\left(\frac{\partial^2 f}{\partial x_p \partial x_q}\right)^2\right\}dx_1 dx_2 dx_3.$$

The solution can be expressed in terms of radial basis functions which parameterise the relationship between points x and y in $\mathbb{R}^3$ as

$$y_d(x) = \sum_{j=1}^{J}\phi\left(||x - x_j||\right)\beta_{jd},$$

where β_{jd} are parameters and d denotes the three dimensions of $\mathbb{R}^3$. Fitting this functional form to the mapping from the observed locations in X to those in Y requires $Y = S\beta_1$, where S is a $J \times J$ matrix, with $S_{ij} = \phi\left(||x_i - x_j||\right)$, and β_1 is a $J \times 3$ matrix whose (j, d)th element is β_{jd}.

It is helpful to separate the mapping into affine and non-affine components, with the former capturing the linear part of the transformation, including possibly different scalings in different co-ordinate directions (shear), and the latter describing non-linear bending. If Q denotes the matrix $(1_J \ X)$, where 1_J is a column vector of 1's, then the transformation can be written in the multivariate form

$$Y = S\beta_1 + Q\beta_2,$$

where β_2 is a 4×3 matrix filled with parameters. This system is now overparametrised, with $(J + 4) \times 3$ parameters but only $J \times 3$ defining equations. This can easily be resolved by adopting suitable constraints, for example through the extended system

$$\begin{pmatrix} Y \\ 0 \end{pmatrix} = \begin{pmatrix} S & Q \\ Q^T & 0 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix},$$ (4)

where the 0 entries indicate matrices filled with 0's of the dimensionality required by the context. These constraints require the sum of the entries of each column of β_1 to be 0 and the sum weighted by the co-ordinates of each dimension of X also to be 0. By applying constraints to the affine component, the interpretation of the non-affine component is left undisturbed.

The system of Equation (4) can be written in the condensed form $Y_e = X_e \beta$, with obvious definitions of X_e and Y_e. If the matrix S is invertible then so is X_e and, after some standard matrix manipulations, the solutions emerge as

$$B_e = \left(S^{-1} - S^{-1} Q \left(Q^T S^{-1} Q \right)^{-1} Q^T S^{-1} \right),$$
$$\beta_1 = B_e Y,$$
$$\beta_2 = \left(Q^T S^{-1} Q \right)^{-1} Q^T S^{-1} Y.$$

When the *bending energy* matrix B_e is post-multiplied by X, this generates the coefficients of the non-affine part of the transformation. The bending energy itself can be expressed as $\mathrm{tr}\ \left\{ Y^T B_e Y \right\}$. Finally, the optimal radial basis function is shown simply to be $\phi(z) = -\frac{1}{8\pi} z$.