A HIERARCHICAL CURVE-BASED APPROACH TO THE ANALYSIS OF MANIFOLD DATA

BY LIBERTY VITTERT[*], ADRIAN W. BOWMAN[*,1] AND STANISLAV KATINA[†,1]

University of Glasgow[*] and Masaryk University[†]

One of the data structures generated by medical imaging technology is high resolution point clouds representing anatomical surfaces. Stereophotogrammetry and laser scanning are two widely available sources of this kind of data. A standardised surface representation is required to provide a meaningful correspondence across different images as a basis for statistical analysis. Point locations with anatomical definitions, referred to as landmarks, have been the traditional approach. Landmarks can also be taken as the starting point for more general surface representations, often using templates which are warped on to an observed surface by matching landmark positions and subsequent local adjustment of the surface.

The aim of the present paper is to provide a new approach which places anatomical curves at the heart of the surface representation and its analysis. Curves provide intermediate structures which capture the principal features of the manifold (surface) of interest through its ridges and valleys. As landmarks are often available these are used as anchoring points, but surface curvature information is the principal guide in estimating the curve locations. The surface patches between these curves are relatively flat and can be represented in a standardised manner by appropriate surface transects to give a complete surface model.

This new approach does not require the use of a template, reference sample or any external information to guide the method and, when compared with a surface based approach, the estimation of curves is shown to have improved performance. In addition, examples involving applications to mussel shells and human faces show that the analysis of curve information can deliver more targeted and effective insight than the use of full surface information.

1. Introduction. One of the interesting types of data generated by medical imaging is in the form of anatomical surfaces. These can arise through the thresholding of three-dimensional (3D) voxel data, for example, to identify the transition between soft tissue and bone. They can also be generated through techniques such as laser scanning or stereophotogrammetry, where surface locations are measured directly through optical methods. Each observational unit consists of a 3D point cloud which provides a discrete, and to some extent noisy, representation of the target surface or manifold.

An immediate issue in the analysis of manifold data is the need to create standardised (*homologous*) representations which have anatomical and geometrical correspondence across images. A traditional approach to this has been through landmarks which identify key point locations; see, for example, Farkas (1994) for a discussion of this in the context of the human face. There is, necessarily, substantial loss of information in reducing the representation to a small number of points, however well chosen, but the wide availability of statistical tools for the analysis of this type of shape information has given landmarks a central role in the study of shape. Dryden and Mardia (2016) give an authoritative account of this topic.

At the other end of the spectrum, a variety of approaches can be taken to the construction of full surface representations. A common approach employs a surface template, which is deformed or warped, so that landmark positions on the template match exactly those on the manifold of interest. Paulsen and Hilger (2003) used thin plate spline warping of landmarks with subsequent adjustment through Markov random field methods, while Hammond et al. (2004) developed a *dense point correspondence* model where standardised nonlandmark locations on the image of interest are identified as the closest points to the corresponding positions on the warped template. In a similar manner, Mao et al. (2006) used local surface curvature to guide further deformation of a template warped by landmarks so as to increase the quality of the match with the manifold in geometrical terms.

Many methods for more general comparison of surfaces have been proposed. Davies, Cootes and Taylor (2001) and Meier and Fisher (2002) described how this can be performed for 2D outlines or 3D surfaces using piecewise linear or smooth functions constructed from appropriate bases. Srivastava et al. (2009) used a Riemannian framework to produce a coordinate system which allows both deformation and comparison using a single elastic metric. A Darcyan curvilinear coordinate system, based on the geodesic distance function from a fixed reference point, allows a facial surface to be represented as an indexed collection of level curves. Rustamov et al. (2013) took a different approach by deriving surface operators which are able to describe the differences between two shapes. In particular, this allows effective interpolation between very different shapes or between very different poses of the same intrinsic shape. Raviv et al. (2014) and Raviv and Kimmel (2015) introduce an equi-affine metric, which creates a form of geometric invariance between nonrigid shapes, while Raviv, Bayro-Corrochano and Raskar (2017) uses a similar approach which allows an average of a collection of surfaces, as well as interpolations and extrapolations of surface sequences, to be computed.

The present paper tackles the problem from a new perspective by focusing attention on anatomical curves which identify ridges and valleys. These are intermediate structures which provide a richer description of manifold data than landmarks but which, as inherently one-dimensional objects, provide simpler representations than a full surface approach. Ridge and valley curves often capture many of the key features of a manifold, as discussed by Koenderink (1990) and many other authors.

For example, in the context of the human face a large proportion of the information on shape is captured in the locations of the ridges and valleys, such as the nose profile or where closed lips meet. Curves may incorporate anatomical landmarks; indeed, Katina et al. (2016) argue that the definitions of anatomical landmarks should, wherever possible, be based on geometrical properties, particularly through the crossing points of anatomical curves or the locations with maximum geodesic curvature. Full surface representations can be approached by relatively straightforward in-filling of areas between curves as the curvature of these intervening regions will be relatively low.

There are therefore strong arguments that ridge and valley curves are the key features to target in constructing representations of shape for manifolds. Indeed, it is one of the aims of the paper to assess the importance of the role of curves in the landmark-curve-surface hierarchy of shape representation. This will be done under the assumption that the curves to be used are present, and homologous, across all the individuals in the study. This corresponds to the standard assumption underlying the analysis of landmarks. Removal of this assumption presents a very interesting but much more challenging situation where different individuals are represented by different sets of curves. However, in the anatomical settings to be discussed in the paper and in many other applications, the assumption of a standard and homologous set of curves is a reasonable one; indeed, this is the basis of the curve definitions proposed by Katina et al. (2016).

In general, the estimation of the locations of ridge and valley curves on a manifold is a challenging problem in the absence of prior information on the nature of the curves to be located. One broad approach is to exploit the definition of ridge or valley points in terms of locally extreme curvature and then to combine candidate points into a curve. Examples of this include Pennec, Ayache and Thirion (2000), Ohtake, Belyaev and Seidel (2004) and Stylianou and Farin (2004) who develop *marching lines* and other forms of algorithm for tracing ridge lines through candidate locations. Che et al. (2011) took a similar approach but assumed the surface of interest to have a smooth parametric form rather than a discrete representation. Bowman et al. (2015) used repeated scans along surface transects to locate points of discontinuity in first derivatives, or sharp change in direction, using statistical methods. Again, ridge and valley curves were created by combining candidate points, although the principal curve (Hastie and Stuetzle (1989)) methods employed do not guarantee a curve which lies on the manifold.

The problem of ridge and valley curve estimation is tackled in the present paper by a new approach which uses basic information on surface curvature to identify regions of interest, transforms these locally to a 2D space and exploits the expected smoothness of the curves to estimate these directly in this new domain. This exploits well-known statistical tools, such as smoothing techniques, but in a formulation which allows these ideas to be used in the novel setting of surface data.

The paper reviews elementary ideas of differential geometry in Section 2 to provide some basic tools for characterising manifolds. Methods for the estimation of ridge and valley curves are described in detail in Section 3 by using the human face as an important example. This leads to the construction of a full surface model in Section 3.4. The properties of the method are evaluated computationally and compared with an existing approach in a simulation study in Section 4. The methods developed are used to study the effects of climate change on the shapes of mussel shells and to quantify human facial sexual dimorphism in Section 5. In both of these examples, analysis based on curve data is found to deliver more targeted and informative insight than the use of full surface information. Some final discussion is given in Section 6.

2. Tools from differential geometry. The local shape at a 3D location $m = (x, y, z)$ on a differentiable manifold $\mathcal{M}$ can be characterised through the quadratic surface

$$(2.1) \qquad z = \frac{1}{2}(\kappa_1 x^2 + \kappa_2 y^2),$$

where z lies in the normal direction to the surface and the orthogonal axes x and y lie on the tangent plane, associated with the directions of maximum (κ_1) and minimum (κ_2) 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 *principal curvatures*, κ_1 and κ_2, and their associated *principal directions*, d_1 and d_2, provide the essential information for characterising curvature across the manifold.

An observation of a manifold consists of M 3D points, each of which is a noise-corrupted version of a point m_i on the manifold $\mathcal{M}$. Specifically, $\mathcal{O} = \{o_i = (x_i, y_i, z_i); i = 1, \ldots, M\}$, where each element can be written as $o_i = m_i + \varepsilon_i$, and ε_i describes the error term associated with the rather complex process of stereophotogrammetric reconstruction. An associated 3D surface triangulation can be constructed from $\mathcal{O}$ through generalisations of the Delaunay approach; see Cohen-Steiner, De Verdiere and Yvinec (2002) or de Berg et al. (2000) for examples. In order to estimate the principal curvatures at an observed point m, a surface normal direction can be constructed as the average of the normal vectors associated with the set of triangles which contain m; see Koenderink and van Doorn (1992). By fitting a quadratic model of the form (2.1) to a local neighbourhood of points through ordinary least squares with arbitrary orthogonal axes in the tangent plane, the principal curvatures and directions can be estimated through the eigen decomposition of the Weingarten matrix. Goldfeather and Interrante (2004) provide all of the details. An important choice is the size of the neighbourhood to which the quadratic surface is fitted. In general, this needs to be adapted to the characteristics of the manifolds being studied. Working with human faces, which at a broad level have very strong shape correspondences across people, allows the effects of different

neighbourhood size to be explored. This resulted in the recommendation of using connected triangulation points within a radius of 1 cm of the point of interest.

There are many useful ways in which the principal curvatures can be condensed into a single number to express particular curvature properties. For example, *Gaussian curvature*, defined as $\kappa_1 \kappa_2$, is a very commonly used measure where positive values correspond to peaks, wells, ridges or valleys, while negative values correspond to saddle points. A particularly useful summary in the present setting is the *shape index* (Koenderink and van Doorn (1992)) which gives a helpful characterisation of the type of curvature at any point. This is defined as

$$(2.2) \qquad S(m) = \frac{2}{\pi} \tan^{-1} \left(\frac{\kappa_2(m) + \kappa_1(m)}{\kappa_2(m) - \kappa_1(m)} \right),$$

where the notation emphasises that the principal curvatures vary with the point of interest m. Values of S close to -1 indicate a 'spherical cup' where both principal curvatures are positive. As S increases, the corresponding surface shape bends smoothly through 'trough' and 'rut' and reaches 'saddle' shapes around 0. The process is reversed as S increases and the associated surfaces move through 'ridge' and 'dome' to arrive at a 'spherical cap' when S reaches 1. This helpful typology of surface shape is illustrated in the upper plot of Figure 1 which is modelled on
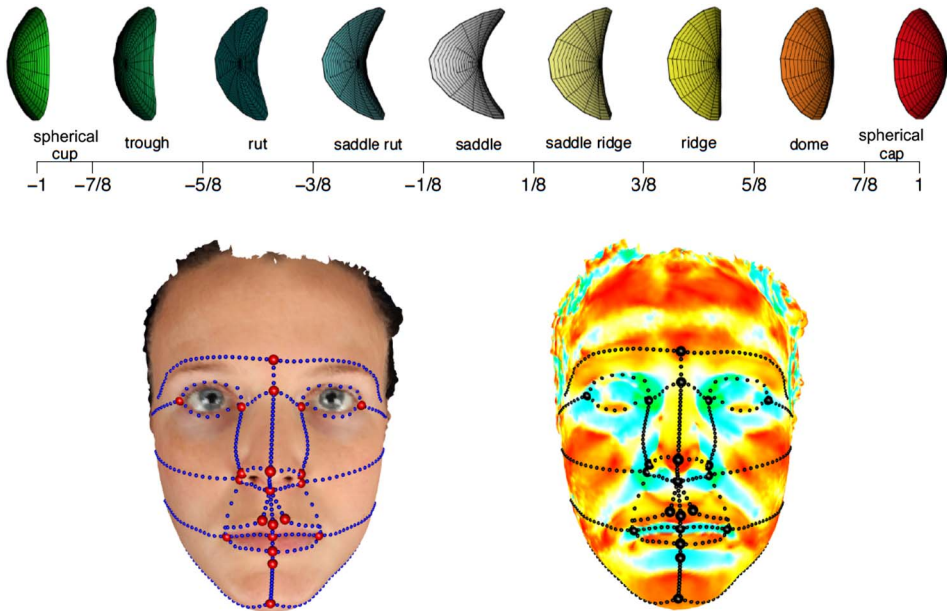

FIG. 1. *The upper plot illustrates the local surfaces associated with the shape index on the scale from −1 to 1, with colour coding to identify each shape category. The lower images show manually placed anatomical landmarks (23 red points), and anatomical curves (467 black points), on facial surfaces coloured by natural texture (left) and shape index (right).*

a similar figure in Koenderink and van Doorn (1992) and uses the same verbal descriptors and colour coding. Notice that the shape index S is a function of the ratio of κ_1 and κ_2, so that it describes the type of local curvature but not its strength.

The lower images of Figure 1 show a human face with the right-hand image coloured by shape index. This highlights key features of the face, such as the ridge of the nose and the valleys and wells around the eyes. The positions of manually placed landmarks and curves align well with key curvature features. This provides encouragement for the use of the shape index and principal curvatures more generally in the estimation of anatomical curves. This topic is developed in the following section.

3. Estimation of ridge and valley curves. Working in a space defined by a manifold requires careful adaptation of the standard procedures used in linear spaces. Patrangenaru and Ellingson (2016) discuss this in detail. To establish notation, it will be useful to refer to a curved path over a manifold as $p(s) = \{x(s), y(s), z(s)\}$, where the functions x, y, z describe the movement of the 3D coordinates as functions of an arc length argument s. The aim is to define paths in the true manifold through curvature and then to estimate these paths by applying the same definitions to the curvature information constructed from the observed manifold. In a repeated sampling framework, or as the image noise is reduced, the convergence of the observed manifold to the true one will ensure that the estimated path also converges to the true one.

Kent, Mardia and West (1996), Stylianou and Farin (2004) and Ohtake, Belyaev and Seidel (2004) discuss the definition of a ridge (or valley) curve. This is most commonly expressed through the characterisation of individual points. A point m with principal curvatures $\kappa_1 > \kappa_2$ and corresponding principal directions d_1, d_2 lies on a valley curve if $\kappa_1 > 0$ and is locally maximal along the direction d_1. The characterisation $D_{d_1}\kappa_1 = 0$, where D_{d_1} denotes the derivative in direction d_1, has consequently promoted methods of estimation based on the determination of zero crossings through numerical and computational techniques. Ridge and valley curves are then estimated through the collection of identified points.

A novel alternative approach is developed here by considering curves from the start. Denote, by $p(s)$, $s \in [0, L]$, a valley curve on a surface of interest. As each point on $p(s)$ satisfies the local optimality condition described above, the integral $\int \kappa_1(p(s))\,ds$ is also locally maximal in the sense that minor perturbations of $p(s)$ must necessarily reduce the contributions made to the integral. Of course, the value of the integral can be increased by allowing the length of $p(s)$ to increase, for example, permitting rapid oscillations which accumulate lower curvatures over longer paths. The optimality of $\int \kappa_1(p(s))\,ds$ is therefore over curves of fixed length. An appropriate estimation strategy then seeks to maximise $\int \kappa_1(p(s))\,ds$ over curve paths p, subject to suitable constraints on path length or smoothness. To allow the simultaneous treatment of ridge curves, where the strongest principal curvature is negative, the notation $\int \nu(p(s))\,ds$ will be used where $\nu(p(s)) = \kappa_1(p(s))$ for valleys and $\nu(s) = -\kappa_2(p(s))$ for ridges.

FIG. 2.   *The upper images show a* 3D *manifold of lips coloured by texture* (*left*), *shape index* (*middle*) *and maximum curvature values* (*right*). *The lower images show three plane cuts to determine the linear path which maximises the length-standardised integral of maximum curvature.*

3.1. *A linear reference path.*   The optimisation of integrated maximal curvature is greatly assisted by the assumption of smoothness. The strategy described below is to identify a local region of interest, use shape index to identify the location of relevant surface shapes and then identify an initial reference curve using an appropriate planar cut. In order to illustrate the issues involved, consider the case of estimating the valley curve where closed lips meet in a human face, as displayed in Figure 2 with two standard landmarks, l_1 and l_2, located at the corners of the mouth. The aim is then to optimise $\int v(p(s))\,ds$ over smooth curves p which have l_1 and l_2 as end-points.

For objects such as human faces, a strong degree of smoothness in ridge and valley curves is to be expected, and this can be quantified in the requirement that the optimal p should lie within the cylinder with axis $(l_2 - l_1)$ and radius r. Empirical investigations with human faces led to the choice $r = \|l_2 - l_1\|/2$ as an effective method of identifying the local region of interest. The top-left images of Figure 2 give an example of the localised region for the lips.

The shape index S over the manifold, illustrated in the top middle image of Figure 2, can now be used to identify those points within the local manifold which indicate valley behaviour. In order to be inclusive at this stage, all points with negative shape index, corresponding to blue-green colour in the standardised scale of Koenderink and van Doorn (1992), are of interest. At locations with positive shape index, v can be assigned the value 0. The top right-hand image of Figure 2 illustrates the result using topographic colour coding.

An initial reference path can now be identified by considering the set of planes containing l_1 and l_2 and indexed by an angle of orientation γ. The lower images of

FIG. 3.   *The left hand image shows lips with anatomical landmarks* (*red*), *a linear reference curve* (*black*), *some illustrative points on the manifold* (*yellow*), *and the corresponding perpendicular distances* (*blue*) *to the closest points on the reference curve* (*green*). *The middle image shows the points from the manifold in the* 2D *space defined by signed distance to the closest point on the reference curve and arc length along the reference curve to the location of the closest point. The colour coding of the points shows the relative strengths of the maximum curvature at each location, on a topographic scale. The right hand image shows a smooth surface representation, with the flexible estimate of the ridge curve in this* 2D *space indicated by the red line.*

Figure 2 illustrate this. For each γ, the intersection of the plane with the manifold describes a path p_γ from l_1 to l_2 with length L_γ. The standardised integral of the maximum curvature along the path p_γ is

$$\left\{ \int \nu\big(p_\gamma(s)\big)\, ds \right\}/L_\gamma \approx \left\{ \sum_{j=2}^{n_p} w_j \nu(p_{\gamma j}) \right\} \Big/ \left\{ \sum_{j=2}^{n_p} w_j \right\},$$

where the right-hand expression shows the discrete approximation based on the path intersection points $\{p_{\gamma j}; j = 1, \ldots, n_{p_\gamma}\}$ on the observed manifold and the weights $w_j = \| p_{\gamma j} - p_{\gamma(j-1)} \|$ measure the distances between successive intersection points. The curve associated with the γ which maximises this expression then 'mops up' as much curvature as possible. The standardisation by curve length L_γ serves to penalise curves which accumulate large amounts of curvature by travelling long distances.

The resulting curve is 'linear' in the sense that it lies in a 2D plane embedded in 3D space, but it describes a curved path over the manifold. The black points in the left-hand image of Figure 3 illustrate the resulting path for the lip valley example.

There are occasions when it is useful to identify a curve across a region which has little strong curvature, such as the cheek from ear to nose landmarks. A linear reference path can easily be constructed simply by minimising over $L_\gamma \approx \sum_{j=2}^{n_p} w_j$ to locate the planar cut with minimum length.

3.2. *A more flexible estimate.*   The problem of finding a path p to maximise $\int \nu(p(s))\, ds$ has been approached through a first approximation which finds the angle γ^* which maximises $\int \nu(p_\gamma(s))\, ds$ over the planar paths p_γ. The use of p_{γ^*} as a reference path allows us to create a helpful coordinate system for further adjustment. Specifically, at any point s, the plane which is orthogonal to p_{γ^*} at

s intersects with the manifold to create a new path $q_s(d)$, where d is a signed distance which indexes the deviations from the reference path in the two possible directions of travel from $p_{\gamma^*}(s)$. This formulation exploits the intrinsically 2D nature of the manifold by creating two new coordinate axes: the signed distance (d) of each point on the manifold from its closest point on the reference path, and the arc length (s) of these closest points along the reference path.

The linear reference path p_{γ^*} can now be adjusted by operating in the (s, d) space to find a function, $r(s)$ say, which describes the magnitude and direction of the orthogonal deviations required to track curvature more effectively. In other words, the function $r(s)$ is chosen to maximise $\int v(p_{\gamma^*}(s) + q_s(r(s)))\, ds$.

This is illustrated in the left-hand image of Figure 3. The distance (d) of a point on the manifold (yellow) to its closest point on the reference path (green) can be measured along the manifold, but simple Euclidean distance is a very good substitute because of the localisation to a region of the manifold very close to the curve of interest. The middle image of Figure 3 uses colour shading to display the values of v_i at the locations (s_i, d_i) for each mesh point ($i = 1, \ldots, n$). This now allows the integral of curvature to be maximised in a standard Euclidean coordinate system.

To address the problem in a continuous rather than a discrete manner, the curvature surface can be conveniently represented in 2D p-spline form (Eilers and Marx (1996)) as

$$(3.1) \qquad v(s, d) = \sum_i \sum_j \hat{\beta}_{ij} \phi_i(s) \phi_j(d),$$

constructed from the product of two one-dimensional cubic b-spline bases $\{\phi_i; i = 1, \ldots, b\}$ and a set of basis weights $\{\hat{\beta}_{ij}; i, j = 1, \ldots, b\}$. The presence of the 'hat' symbol on each $\hat{\beta}_{ij}$ reflects the fact that these have been estimated by fixing the equivalent degrees of the estimator to be 12, which allows a good degree of flexibility for the fitted surface. The aim is to locate the optimal curve, for which a one-dimensional p-spline representation could also be used. However, an even simpler approach is to represent the optimal curve r at a grid of arc length positions $\{s_k; k = 1, \ldots, n_g\}$ as

$$r(s_k) = \alpha_k, \quad k = 1, \ldots, n_g.$$

Following the earlier principle that the integral of curvature along the optimal curve is locally maximal, the aim is to identify $r(s)$ to maximise

$$\int v(s, r(s))\, ds.$$

By adopting a discrete approximation and by exploiting the assumption of smoothness in the optimal curve through a penalty function, this translates into the problem of identifying the $\boldsymbol{\alpha}$ (the vector of α_k values), which maximises the penalised

discrete integral

$$G = \frac{1}{n_g} \sum_k \nu(s_k, \alpha_k) - \lambda \boldsymbol{\alpha}^T P \boldsymbol{\alpha}$$

(3.2)

$$= \frac{1}{n_g} \sum_k \sum_i \sum_j \hat{\beta}_{ij} \phi_i(s_k) \phi_j(\alpha_k) - \lambda \boldsymbol{\alpha}^T P \boldsymbol{\alpha},$$

where λ denotes a penalty parameter and $P = D^T D$ denotes the penalty matrix constructed from the matrix D, which generates second-order differences of the elements of α as $D\alpha$. As the expression G is not linear in α, Newton's method provides a suitable solution. This requires the evaluation of the derivatives

$$\frac{\partial G}{\partial \alpha_k} = \sum_i \sum_j \hat{\beta}_{ij} \phi_i(s_k) \phi'_j(\alpha_k) - 2\lambda(P\alpha)_k,$$

$$\frac{\partial^2 G}{\partial \alpha_k^2} = \sum_i \sum_j \hat{\beta}_{ij} \phi_i(s_k) \phi''_j(\alpha_k) - 2\lambda(P)_{k,k},$$

$$\frac{\partial^2 G}{\partial \alpha_k \partial \alpha_l} = -2\lambda(P)_{k,l} \quad (k \neq l),$$

where the notation $(P\alpha)_k$ denotes the kth element of the vector $P\alpha$ and $(P)_{k,l}$ denotes the (k, l)th element of the matrix P. B-splines have the very convenient property that their derivatives are scaled b-splines of lower order. Specifically,

$$\phi'_{i,a}(x) = (\phi_{i,a-1}(x) - \phi_{i+1,a-1}(x))/b,$$

where the second subscript on ϕ denotes the order of the b-spline function and b denotes the distance between the (regularly spaced) knots. Second derivatives can then be computed by a further application of this formula.

The anchoring points correspond to $\alpha_1 = 0$ and $\alpha_{n_k} = 0$, so these two elements are fixed. From a starting point, which sets all $\alpha_k = 0$, Newton's method then employs the iterations

$$\boldsymbol{\alpha}^{(m+1)} = \alpha^{(m)} - H^{-1}(\alpha^{(m)}) f(\alpha^{(m)}),$$

where f denotes the vector of first derivatives and H denotes the matrix of second derivatives. The iterations converge very quickly.

After experimentation on a variety of images, the penalty parameter was set to $\lambda = 0.5$ as this achieves a good compromise between fidelity and smoothness for the relatively simple curves in human faces. The right-hand images of Figure 4 show the end result of this process on the illustrative example of a lip valley curve. The resulting curve has a flexibility, which is modest but sufficient to adapt to the curvature information, which suggests that the true valley curve deviates a little from the earlier linear reference path.

Three technical details of the estimation method are discussed in the Supplementary Material (Vittert, Bowman and Katina (2019a)).
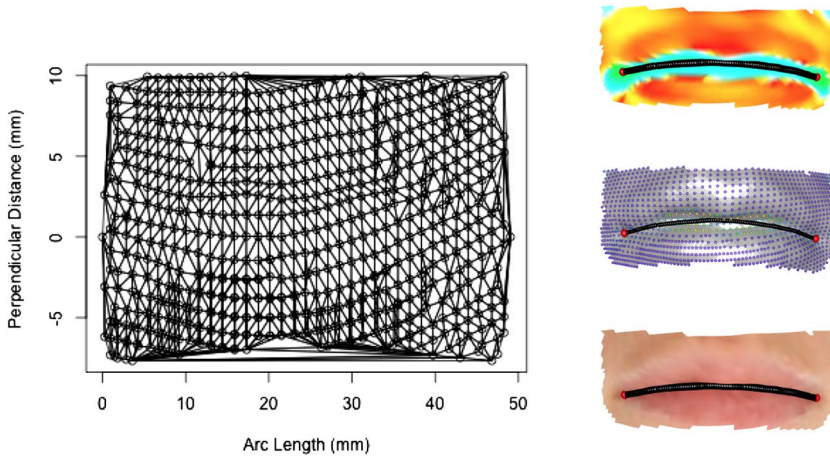
FIG. 4. *The left hand plot shows a Delaunay triangulation, which forms the basis of barycentric interpolation. The right hand images show the final estimate of the lip valley curve, on surfaces coloured by shape index (see Figure 1 for colour scale), maximum curvature and natural texture.*

3.3. *Interpolation back to 3D space.* The final step is to transfer the estimate of the valley curve from the 2D (s, d) domain back to 3D space. The existence of a triangulation on the observed manifold means that barycentric interpolation can provide a simple solution. There is a 1–1 correspondence between the observed points o_i on the manifold and the corresponding observed points (s_i, d_i) in the 2D space. As the triangulation can be expressed through sets of triplets (j_1, j_2, j_3), which give the indices of the observed points $(o_{j_1}, o_{j_2}, o_{j_3})$ on the manifold which form the vertices of the jth triangle, the triangulation can therefore also be expressed in the 2D space by forming the triangles from the same sets of indices. The left-hand panel of Figure 4 illustrates this. For any point (s, d) in the 2D space, for example, a point on the estimated valley curve, the first step is to find the enclosing triangle, defined by the indices (j_1, j_2, j_3). Barycentric coordinates allow (s, d) to be written as $(\sum_{k=1}^{3} \lambda_k s_{j_k}, \sum_{k=1}^{3} \lambda_k d_{j_k})$ for a unique set of weights $(\lambda_1, \lambda_2, \lambda_3)$, where $0 \leq \lambda_1, \lambda_2, \lambda_3 \leq 1$. These interpolation weights for the position of (s, d) can now be used to interpolate other values. For example, the associated 3D coordinates at (s, d) can be interpolated from the 3D coordinates of $(o_{j_1}, o_{j_2}, o_{j_3})$. Specifically, the x coordinate can be interpolated as $\sum_{k=1}^{3} \lambda_k x_{j_k}$ with the y and z coordinates following suit. The linear form of the interpolation ensures that the computed location lies on the 3D triangulated surface. Meyer et al. (2002) provide the details.

The images on the right-hand side of Figure 4 display the resulting estimate of the lip valley curve on surfaces, which are coloured by shape index and maximum curvature, as well as by natural texture, in order to highlight the information, which has been used in the construction of the estimate. This process can be repeated

across all the curves of the face displayed in Figure 1. In fact, Figure 1 shows the curves estimated by the methods described in this section.

3.4. *Construction of a facial surface model.* A set of anatomical curves across the face creates a structure, which already captures a considerable amount of the information on shape. A full representation of the facial surface can then be produced by characterising the patches which have these curves as boundaries. As the construction of the curves targets locations of high curvature, these patches will necessarily have lower curvature but they may still express important biological signals. Discrete representation of the anatomical curves, using suitably spaced points along each, then allows a description through linear transects between these boundary points, using plane cuts minimised over length as described in Section 3.1.

However, care must be taken in the construction of a discrete representation. This is an issue which often occurs in functional data analysis, for example, where growth curves exhibit peaks at different times across individuals. Naive averaging across curves at each time point blurs the signals displayed in each individual curve. Kneip and Ramsay (2008) show how 'structural averaging' can be achieved by employing a transformation of the timescale for each curve to align on key features. In the context of shape, corresponding issues arise, and Bookstein (1997) describes a related idea of 'sliding landmarks' where the positions of points are successively adjusted along the curve to minimise bending energy with respect to a template curve. This is a widely used technique in geometric morphometrics, and Rohr (2001) discusses the connection with principal curvatures. In general, the problem is considerably reduced in settings where there are sufficient numbers of landmarks to allow the intervening curves to have relatively simple, smooth shapes. This is the case in the examples discussed in Section 5 below.

4. Evaluation of the performance of curve estimation. The key step in the construction of the facial model described above is the estimation of the ridge and valley curves. In order to evaluate the proposed method, a simulation study was undertaken to compare estimated curves with known ridge locations. In addition, estimated curves were compared with curves identified manually by a trained observer on a set of real images.

4.1. *Simulation study.* Simulated surfaces were generated to mimic the shape of human lips over the planar domain $(x, y) \in [-1, 1] \times [-1, 1]$. The equations below use simple quadratics to define the locations of a valley curve v and two adjacent ridge curves r_1, r_2. The function $z(x, y)$ then creates the surface as a superposition of two ridges with Gaussian cross-sections, placed along the ridge curve locations. The numerical coefficients of the curves and ridges have been selected to create a well-defined valley between the two ridges. The parameter c
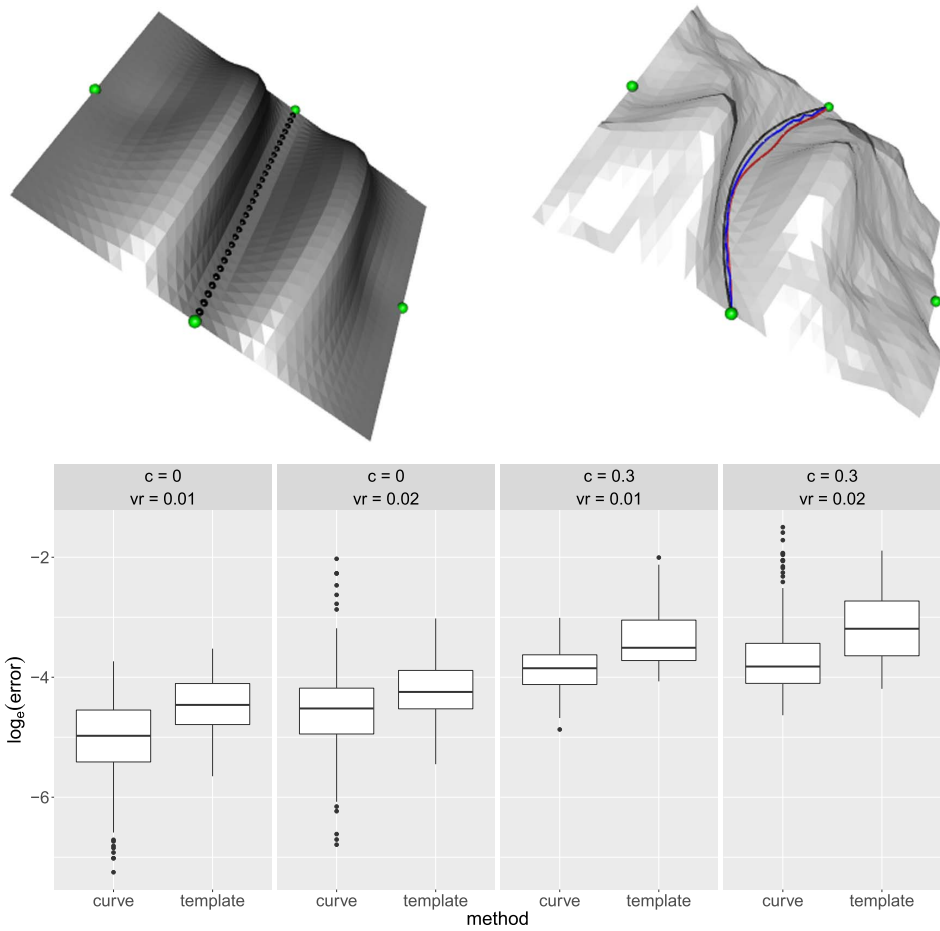
FIG. 5. *The upper left hand image shows a surface template. The right hand image shows an example of a simulated surface, as described in the text, with parameters $c = 0.3$ (midline curvature), $h = 0.6$ (spatial correlation range) and $vr = 0.02$ (variance). The true midline curve is shown in black, while the estimates produced by the direct and template methods are shown in blue and red respectively. In both of the images, the landmarks are shown in green. The lower plot show the performance of the direct and template methods in terms of the average distance between the true midline curve and the estimates, for a variety of different simulation settings.*

controls the curvature of the valley curve. The top left-hand image of Figure 5 illustrates the resulting surface over a regular $n \times n$ grid with $c = 0$.

    Surface :

$$v(x) = c(1 - x^2),$$
$$r_1(x) = v(x) + 0.01(1 - x^2),$$

$$r_2(x) = v(x) - 0.01(1 - x^2),$$

$$z(x, y) = 5\big(e^{(-\frac{1}{2}(\frac{y-r_1(x)}{0.25})^2)} + e^{(-\frac{1}{2}(\frac{y-r_2(x)}{0.25})^2)}\big)(1 - x^2)(y - v(x))^2.$$

The $(1 - x^2)$ terms in r_1, r_2 and $z(x, y)$ act as weight functions, which bring the ridges and the surface down to 0 at the left- and right-hand edges, ensuring that the landmarks $l_1 = (-1, 0, 0)^T$ and $l_2 = (1, 0, 0)^T$ lie on the surface.

In order to create a distribution of shapes, the surface $(x, y, z(x, y))$ was perturbed by the addition of a Gaussian spatial process $\varepsilon(x, y)$, using a Matérn covariance function with shape parameter 1.5, variance parameter vr and correlation range parameter h. A weight function w_v was used to ensure that the shape of the surface near the midline curve, and hence its definition as a valley, remained undisturbed. The template method described below requires further landmarks, and these were defined as $l_3 = (0, 1, 0)^T$ and $l_4 = (0, -1, 0)^T$. Further weight functions, w_3 and w_4, were therefore used to ensure that these additional landmarks lay on the surface after the addition of the spatial process. The weight functions and the resulting surface definition are shown below.

Simulation :

$$w_v(x, y) = \frac{e^{(40(\delta_v(x,y)-0.15))}}{(1 + e^{(40(\delta_v(x,y)-0.15)))}},$$

$$w_3(x, y) = 1 - e^{-\frac{1}{2}(\frac{\|(x,y,z(x,y))^T - l_3\|}{0.2})^2},$$

$$w_4(x, y) = 1 - e^{-\frac{1}{2}(\frac{\|(x,y,z(x,y))^T - l_4\|}{0.2})^2},$$

$$f(x, y) = z(x, y)w(3(x, y)w_4(x, y) + \varepsilon(x, y)w_v(x, y)w_3(x, y)w_4(x, y).$$

An example of a surface, with the midline curvature parameter set to $c = 0.3$, is displayed in the top right-hand image of Figure 5.

To provide a comparison with a standard approach, the surface template method described by Mao et al. (2006) was also used. In the absence of packaged software to implement this routinely, an analogue of the method was constructed in a manner which allows as close correspondence as possible between the curve and surface approaches. The first step is to construct a template which encapsulates the main characteristics of the surfaces to be modelled, as far as possible. The surface described at the beginning of Section 4.1 above was used, with $c = 0$, as displayed in the top left image of Figure 5. The template is then 'warped' onto the surface image of interest, using thin plate splines to match the two sets of landmarks exactly while smoothly transforming the remainder of the template surface. This is a standard technique in statistical shape analysis described, for example, by Dryden and Mardia (2016). Local deformation is then carried out by considering how the location of each point t on the template should be further adjusted to match the shape characteristics of the image more effectively. Specifically, a measure of the dissimilarity of t and a point i on the image is computed

as $\|t - i\|/r + \cos^{-1}(n_t^T n_i)/\pi + |S_t - S_i|$, where S_t, S_i and n_t, n_i denote the shape indices and normal vectors at t and i. All locations i within a distance R of t were considered and the point i_t with minimum dissimilarity identified. (The distance parameters were set as $R = 2$ and $r = 0.4$, following the advice of Mao et al. (2006)). The final step is to ensure that the surface created by moving each t to i_t is smooth and coherent. Mao et al. (2006) use energy minimisation to achieve this, but, in order to achieve maximum comparability with the curve estimation method described in Section 3, 2D p-spline smoothing, as in (3.1), of each coordinate of the adjustment vector $i_t - t$ was performed over the $n \times n$ grid on the $[-1, 1] \times [-1, 1]$ spatial domain of the template. When this process has been completed, the valley curve in the template, represented by the points in a straight line between the landmarks l_1 and l_2, now occupies a new position, and this is taken as the estimate of the valley curve in the image.

In order to allow each method to perform to best effect in the simulations, the mean separation (average distance between 50 equally spaced points along each curve) of the estimated and true valley curves was computed over a variety of settings of the smoothing control parameters. In the template method, this is controlled by df_t, the degrees of freedom (Eilers and Marx (1996)) used in smoothing the adjustment vectors. In the curve method, smoothing is split across the degrees of freedom df_c used in the construction of the 2D surface (3.1) and the penalty parameter λ used in ridge estimation, described in (3.2). The most effective overall performance was achieved by $df_t = 50$, as considerable flexibility is required to match complete surfaces, $df_c = 12$, which confirms the value reported in Section 3.2, and $\lambda = 2000$, which adapts to the scales of the spatial coordinates used in the simulation study.

The two methods were then compared over a variety of parameters controlling the shapes of the simulated surfaces, namely c (shape of the valley curve), vr and h (variance and range of the added spatial process). In each case, 200 simulations were created, using a regular grid of 31×31 points across the surface. Figure 5 shows the mean distances (on a log scale to adjust for skewness) between the true and estimated valley curves in the case of $h = 0.6$. In each pair of boxplots, the direct curve estimate shows superior performance. There are some isolated cases of higher errors in curve estimation, but these are due to reference paths which occasionally locate valley signals in rapidly fluctuating surfaces adjacent to the true valley curve. This occurs only at levels of noise ($vr = 0.2$) which are much higher than those exhibited in the applications later in the paper. In general, the results highlight the benefit of estimating a curve by focussing on the curve explicitly, rather than creating the estimate as a by-product of a more general surface matching method. The absence of any need to create a suitable template is also a significant asset of the curve-based approach. The guarantee that the curve lies on the image surface, as discussed in Section 3, is also a major advantage which is not enjoyed by the template matching approach.

Further details of bias calculations and simulation results are reported in the Supplementary Material (Vittert, Bowman and Katina (2019b, 2019c and 2019d)). In particular, these show that the effects of inaccurate placement of landmarks is modest.

4.2. *Manual comparison.* Manual identification by trained individuals has been regarded as the 'gold standard' for the production of anatomical landmarks. It needs to be recognised that manual methods are also subject to variation and inaccuracy, and that this will undoubtedly increase in the manual identification of curves rather than points. Nonetheless, it is valuable to compare the curves produced by the facial model described above with those identified manually, here using the *Landmark* ©IDAV software. The facial model curves and the manually marked curves were constructed on 55 images and represented by sets of discrete points to allow Procrustes registration of each pair of curves. This gave an overall root mean square distance of 1.70 mm between corresponding points. Differences between the curves were strongest in the lip borders and the nasal bridge. When these regions are omitted, the overall root mean square distance between curves drops to 0.60 mm.

It is instructive to consider the reasons why these two areas show the largest differences. When marking lips, manual observers are likely to be strongly influenced by colour change at the so-called 'vermillion border.' However, it is not always the case that the ridge of the lips coincides exactly with this colour change. At the nasal bridge, there is a tendency for manual observers to mark positions which are slightly 'north' of the ridge curve. It can be very difficult to make a visual judgement on where this ridge lies, with a great deal of examination of the image required at different orientations. In contrast, the facial model employs direct information on surface curvature and is unaffected by perceptual considerations.

5. Applications. Two examples of shape analysis are investigated in this section involving the shapes of mussel shells and human faces. The aim is to illustrate the curve estimation methods described in Section 3 and in particular to investigate the information and insight delivered by the different levels of the landmark-curve-surface hierarchy. Each of these different levels of information can be analysed by methods which exploit their specific characteristics. For example, Cheng, Dryden and Huang (2016) describe various approaches specific to curve data. However, in order to allow simple and even-handed comparison across the different types of information, Principal Component Analysis (PCA) will be used as this has sufficient generality to be applicable to all these forms.

5.1. *Mussel shells and climate change.* Fitzer et al. (2015) describe a study of the potential effects of climate change on the shape of mussel shells, separated into *six* different groups of *four* to assess the affects of growth, over a nine-month period, under increasing environmental pressure, based on ocean acidification and
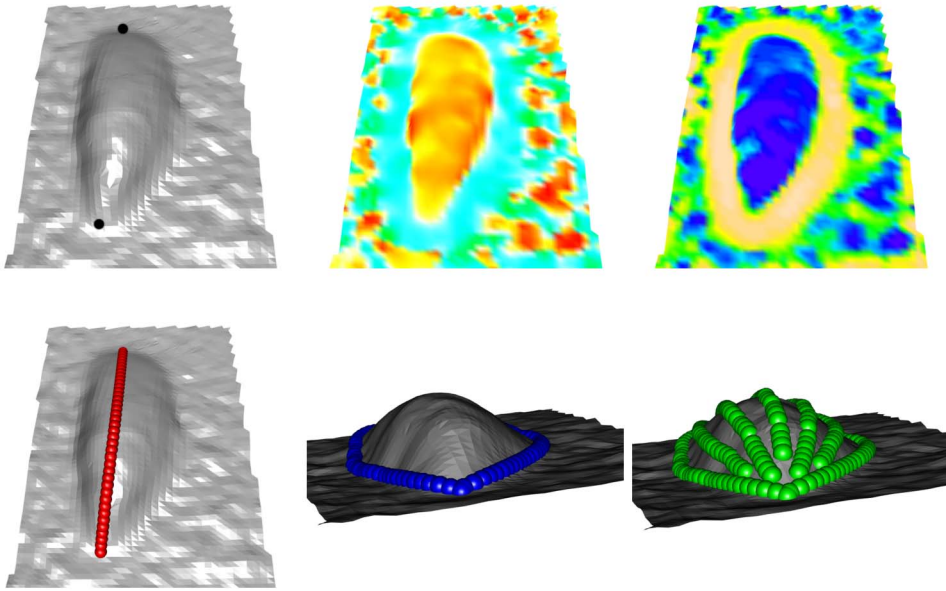
FIG. 6.   *The upper images show a mussel shell* (*left*) *coloured by shape index* (*middle*), *and maximum curvature values* (*right*). *The lower images illustrate the central ridge* (*left*) *of the shell*, *the perimeter* (*middle*) *and a surface representation based on interpolation between these* (*right*).

temperature. The top images of Figure 6 show a 3D image of a mussel shell placed on a flat surface, with two manually marked, anatomical landmarks, and versions of the same image coloured by shape index and by principal curvature κ_1. This last image highlights the perimeter curve. The central ridge of the mussel shell, and its perimeter, were identified using the ridge and valley curve estimation methods described in Section 3. A full surface model was then created by constructing shortest distance paths between the ridge and perimeter curves. These steps are all illustrated in Figure 6.

PCA was performed on the shape representations at each level of the landmark-curve-surface hierarchy, but in this case the number of landmarks is too small to provide any useful information. For the curves and surfaces, the first *four* principal components (PCs) were considered as these explain over 90% of the variability present. Table 1 shows for each component the evidence for change in shape with environmental pressure, expressed in the form of p-values associated with a test of no effect based on nonparametric regression; see Bowman and Azzalini (1997) for details. No significant effects are apparent in the surface representations, but the second and third PCs show evidence of change in the curve shapes, particularly in the perimeter. The nature of the shape changes in the perimeter is indicated in the lower and upper right-hand images, respectively, of Figure 7 through the curves associated with ± 3 standard deviations from the mean in this component direction. As environmental pressure increases, the mussel perimeter tends to widen

TABLE 1
P-values associated with nonparametric regression with environmental pressures against shape change (*PCs*)

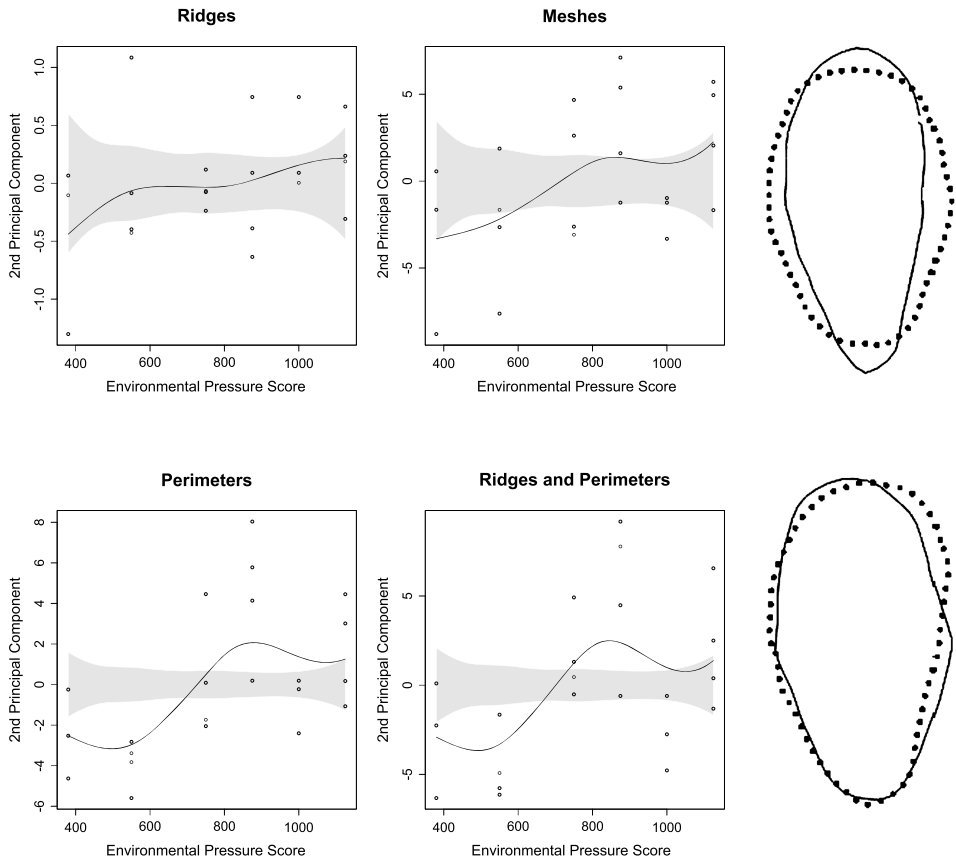| PC	Ridge	Perimeter	Ridge and Perimeter	Surface
1	0.997	0.966	0.986	0.992
2	0.470	0.024	0.030	0.156
3	0.974	0.020	0.260	0.536
4	0.414	0.831	0.994	0.995

FIG. 7. *The four left hand plots relate the scores for PC* 2 *to environmental pressure, using information on the perimeter curve, the ridge curve, both of these curves, and a full surface representation. The two right hand plots show the perimeter shape changes associated with PCs* 2 (*upper*) *and* 3 (*lower*), *represented through the curves associated with* +3 (*dotted*) *and* −3 (*solid*) *standard deviations from the mean of the corresponding scores.*

with a more rounded shape. The third PC shows a change in the orientation of the 'pointed' end of the mussel tip.

The absence of evidence of shape change, when a full surface model is used, highlights the potential benefit of using the focussed information provided by curves, which are unaffected by the variability, which may be present in the more detailed description of the intervening surface shape.

The results reported here are broadly consistent with those reported by Fitzer et al. (2015). Additional seawater measurements were used in the earlier paper and a slightly larger set of data was used in this paper, as a result of improved performance in the curve estimation methods. (Reference to surface analysis in the earlier paper essentially corresponds to a combination of ridge and perimeter curves.)

5.2. *Sexual dimorphism of the human face.*   The differences in facial shape between human males and females have been studied from a variety of perspectives. See, for example, Bruce et al. (1993) and Armann and Balthoff (2012) for one example from visual perception. The manner in which shape is quantified clearly has a very big influence on the information which is then available for analysis. In this area, shape has traditionally been quantified through distances, angles and ratios computed from landmarks in 2D and 3D images, but the information present in landmarks is necessarily limited. A recent example of a full 3D surface representation of faces is described by Claes et al. (2012) who warp a facial template onto the images and focus particularly on symmetric/asymmetric components of sexual dimorphism. It is of interest to use the curve estimation algorithm reported in Section 3 to investigate the information contained in the different levels of landmark, curve and surface representations in this context.

A sample of 219 facial images captured from volunteers (ages 18–45 years) was used. All analysis reported in this section compares males and females through a permutation approach, with 1000 permutations of sex labels to avoid reliance on distributional assumptions. A larger mean centroid size is expected for males, and this was confirmed by a two sample t-test ($p = 0.011$). Procrustes registration, with scale removed, was used to focus attention on shape free of scale. Further analysis was based on PCs as this provides a standard mechanism to describe the main modes of shape variation in a smaller number of dimensions, as discussed by Dryden and Mardia (2016), for example. The first *six* PCs were used here as this captures around 75% of the shape variation. A Hotelling's T^2 test of differences in mean facial shape between males and females, based on these components for curves, found a significant difference ($p = 0.018$).

Table 2 shows the evidence of sexual dimorphism from individual PCs across all three forms of shape representation. Each row of Figure 8 corresponds to a particular PC with boxplots to show the differences in scores between males and females and facial images to illustrate the nature of the associated shape change. Topographic colours are used to show the movement between the extremes at ±3

TABLE 2
P-values associated with a two-sample comparison of males and females

PC	Landmarks	Curves	Meshes
1	0.390	0.446	0.523
2	0.510	0.028	0.627
3	0.202	0.001	0.000
4	0.200	0.003	0.403
5	0.000	0.000	0.000
6	0.200	0.013	0.185

standard deviations away from the mean score, starting at the end of the component scale dominated by females and moving to the end dominated by males (The topographic scale is shown in millimeters at the bottom of Figure 8.) The distances here are based on curves, but a full facial image has been warped onto the curve positions to highlight the interpretation of these changes more effectively. In the Supplementary Material (Vittert, Bowman and Katina (2019e), Figure 9), the extremes are shown for the other two levels (landmarks and mesh) of hierarchical analysis for PCs $2-5$, clearly showing that the shape change associated with each component is similar across all levels of the hierarchy.

Table 2 shows that PC 5 is clearly significant in all three hierarchical levels. Figure 8 shows that this corresponds to movement in the *gnathion* (lower chin) which is generally lower in males than in females. PC 3 is significant in both the curve and mesh analyses. This corresponds to an increased relative size of the nasal region and a higher periorbital region for males.

Table 2 also shows that PCs 2 and 4 are significant in the curve analysis but are not significant in the mesh analysis. The shape changes in PC 2 correspond to a more prominent jaw and supraorbital region (brow) in males. PC 4 corresponds to a lowering of the brow ridge and strong lowering of the jaw in males. (Notice that the oral region also shows slight forward movement.) In contrast, these shape changes are not expressed as strongly in the components constructed from the full surface model.

Following on from the simple example of the mussel analysis, the curve estimation algorithm described in this paper delivers a clearer analysis than the prevalent conformed mesh methods. Again, we see that the most significant facial information is held in the facial curves and that significant shape changes can be obscured when too much unnecessary information is introduced in a surface representation. The analysis indicates the key role played by the facial curves, where much of the shape information resides, in a similar manner to the effects observed in the mussel shell analysis.
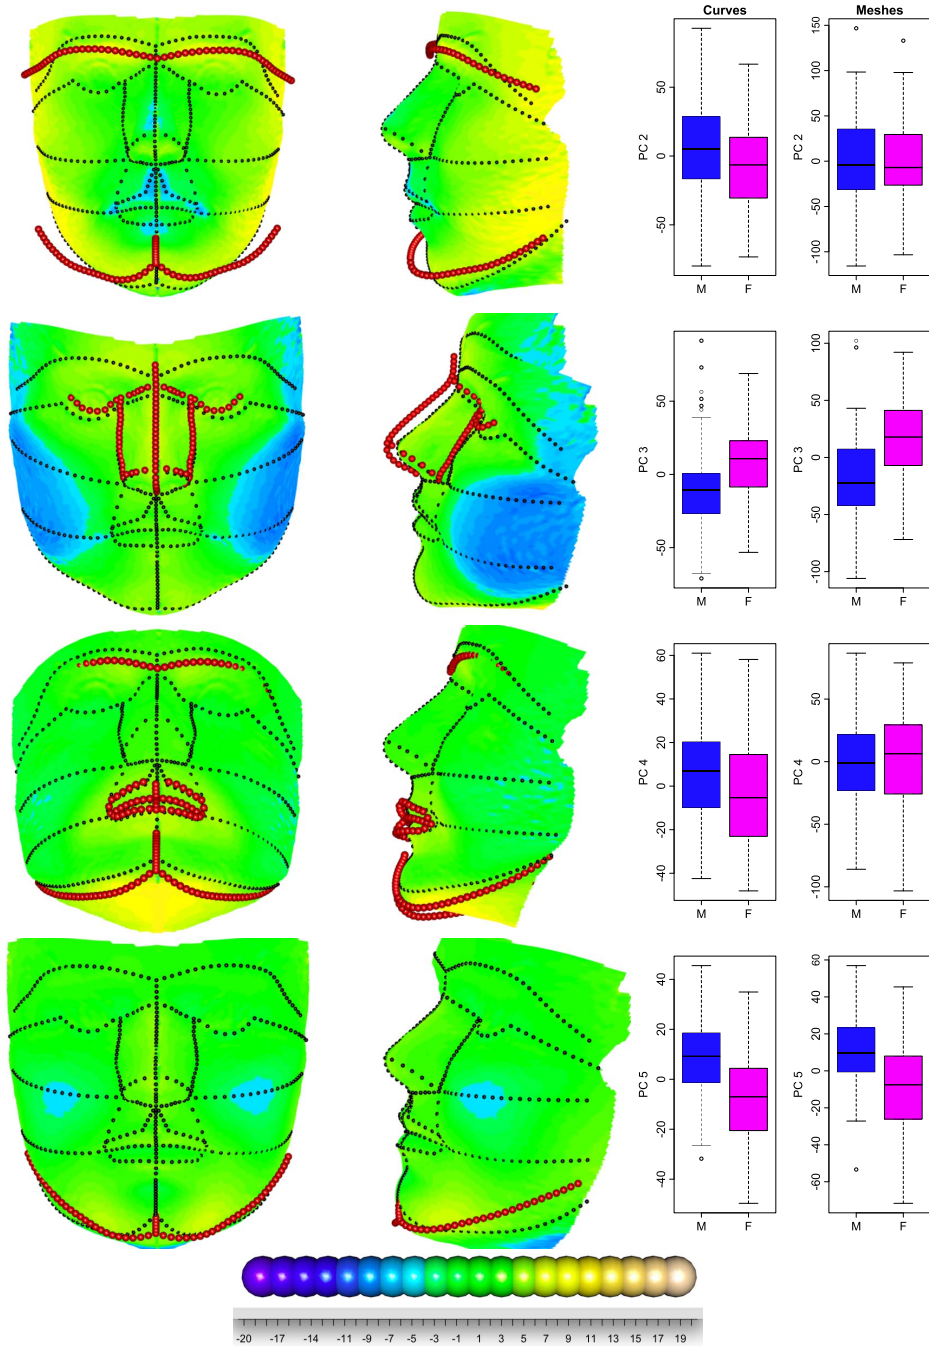
FIG. 8. *Extremes of the curve analysis*: 2*nd*, 3*rd*, 4*th* and 5*th PC respectively, the female face coloured topographically by the shape change to male with the male curves in red and female curves in black. The topographic scale in millimeters is shown at the bottom of the figure.*

6. Discussion. The approach proposed in the paper gives a full description of a 3D manifold through estimation of its ridge and valley curves and subsequent 'in-filling' of the intervening surface patches. The methods used respect the inherently 2D nature of the manifold by operating in relevant 2D spaces. The maximisation of curvature along path integrals is used as the guiding principle for the estimation of the ridge and valley curves, with efficient implementation through p-spline representations. The applications illustrate that the information contained in the curves can inform on shape differences more effectively than representations based simply on landmarks or indeed on the complexity of full surfaces. This endorses the viewpoint that in the point-curve-surface hierarchy of manifold descriptions the curve information plays a particularly important role, achieving a parsimonious but informative descriptor of shape. There are many other application areas, particularly in human and wider biological anatomy, where the analysis of surface shape is central and where we believe that a curve based approach could be beneficial. A reviewer helpfully suggested that brain cortical surfaces would be a fascinating application, although the lack of homologous structure across individuals and the tree structure of the curves involved introduce complexity which would require further methodological development.

Methods based on surface templates are very dependent on the particular template used. Even in the case of human faces, which display strong similarities across different populations, care must be exercised in the choice of template. For example, experimentation with a 'White British' template in modelling facial images from subjects with African origin led to artefacts and distortions on several occasions. The use of curves, which are estimated directly from the image of interest, are free of this complication.

A further benefit is that individual features of surfaces (nose, lips, etc.) can easily be defined by nominating curves as boundaries. Figure 2 illustrates this in the case of the human face where lips, nose and other features are immediately available as subsets of the image.

The analysis here has exploited the presence of traditional landmarks by using these to 'seed' the construction of the ridge and valley curves. The existence of methods of estimating these curves opens up the possibility of informing on the location of landmarks which can be defined in terms of curve crossings or points of extreme geodesic curvature. The combined estimation of curves and landmarks is therefore a natural topic for further research.

SUPPLEMENTARY MATERIAL

Supplementary Material A: Technical details of the ridge and valley curve estimation (DOI: 10.1214/19-AOAS1267SUPPA; .pdf). We provide additional supporting technical details for the original curve and valley estimations.

Supplementary Material B: Simulations to compare curve and surface based methods of curve estimation (DOI: 10.1214/19-AOAS1267SUPPB; .pdf). We provide additional details on the simulations that compare the curve based method presented in this paper and the surface based methods by others'.

Supplementary Material C: Bias in the estimation of surface curvature (DOI: 10.1214/19-AOAS1267SUPPC; .pdf). We document the bias that exists in the estimation of the surface curvatures.

Supplementary Material D: Simulations to assess the influence of landmark placement (DOI: 10.1214/19-AOAS1267SUPPD; .pdf). We provide additional supporting tables that display the change in curve estimations based upon initial landmark placement.

Supplementary Material E: Principal components for sexual dimorphism (DOI: 10.1214/19-AOAS1267SUPPE; .pdf). We provide additional supporting plots that show the further Principal Component information for the sexual dimorphism discussion.

REFERENCES

ARMANN, R. and BALTHOFF, I. (2012). Male and female faces are only perceived categorically when linked to familiar identities. And when in doubt, he is a male. *Vis. Res.* **63** 69–80.

BOOKSTEIN, F. L. (1997). Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape. *Med. Image Anal.* **1** 225–243.

BOWMAN, A. W. and AZZALINI, A. (1997). *Applied Smoothing Techniques for Data Analysis*: *The Kernel Approach with S-Plus Illustrations*. Oxford University Press, Oxford.

BOWMAN, A. W., KATINA, S., SMITH, J. and BROWN, D. (2015). Anatomical curve identification. *Comput. Statist. Data Anal.* **86** 52–64. MR3312737

BRUCE, V., BURTON, A. M., HANNA, E., HEALEY, P., MASON, O., COOMBES, A., FRIGHT, R. and LINNEY, A. (1993). Sex discrimination: How do we tell the difference between male and female faces? *Perception* **23** 131–152.

CHE, W., ZHANG, X., ZHANG, Y.-K., PAUL, J.-C. and XU, B. (2011). Ridge extraction of a smooth 2-manifold surface based on vector field. *Comput. Aided Geom. Design* **28** 215–232. MR2802876

CHENG, W., DRYDEN, I. L. and HUANG, X. (2016). Bayesian registration of functions and curves. *Bayesian Anal.* **11** 447–475. MR3471998

CLAES, P., WALTERS, M., SHRIVER, M. D., PUTS, D., GIBSON, G., CLEMENT, J., BAYNAM, G., VERBEKE, G., VANDERMEULEN, D. et al. (2012). Sexual dimorphism in multiple aspects of 3d facial symmetry and asymmetry defined by spatially dense geometric morphometrics. *J. Anat.* **221** 97–114.

COHEN-STEINER, D., DE VERDIERE, E. C. and YVINEC, M. (2002). Conforming Delaunay triangulations in 3D. In *Proceedings of the Eighteenth Annual Symposium on Computational Geometry* 199–208. ACM. New York.

DAVIES, R. H., COOTES, T. F. and TAYLOR, C. J. (2001). A minimum description length approach to statistical shape modelling. In *Biennial International Conference on Information Processing in Medical Imaging* 50–63. Springer.

DE BERG, M., VAN KREVELD, M., OVERMARS, M. and SCHWARZKOPF, O. (2000). *Computational Geometry*: *Algorithms and Applications*, revised ed. Springer, Berlin. MR1763734

DRYDEN, I. L. and MARDIA, K. V. (2016). *Statistical Shape Analysis*, 2nd ed. Wiley, New York. MR1646114

EILERS, P. H. C. and MARX, B. D. (1996). Flexible smoothing with *B*-splines and penalties. *Statist. Sci.* **11** 89–121. MR1435485

FARKAS, L. (1994). *Anthropometry of the Head and Face*, 2nd ed. Raven Press, New York.

FITZER, S. C., VITTERT, L., BOWMAN, A., KAMENOS, N. A., PHOENIX, V. R. and CUSACK, M. (2015). Ocean acidification and temperature increase impact mussel shell shape and thickness: Problematic for protection? *Ecol. Evol.* **5** 4875–4884.

GOLDFEATHER, J. and INTERRANTE, V. (2004). A novel cubic-order algorithm for approximating principal direction vectors. *ACM Trans. Graph.* **23** 45–63.

HAMMOND, P., HUTTON, T. J., ALLANSON, J. E., CAMPBELL, L. E., HENNEKAM, R., HOLDEN, S., PATTON, M. A., SHAW, A., TEMPLE, I. K. et al. (2004). 3d analysis of facial morphology. *Am. J. Med. Genet.*, *Part A* **126** 339–348.

HASTIE, T. and STUETZLE, W. (1989). Principal curves. *J. Amer. Statist. Assoc.* **84** 502–516. MR1010339

KATINA, S., MCNEIL, K., AYOUB, A., GUILFOYLE, B., KHAMBAY, B., SIEBERT, P., SUKNO, F., ROJAS, M., VITTERT, L. et al. (2016). The definitions of three-dimensional landmarks on the human face: An interdisciplinary view. *J. Anat.* **228** 355–365.

KENT, J. T., MARDIA, K. V. and WEST, J. (1996). Ridge curves and shape analysis. In *BMVC Proceedings* 1996 1–10.

KNEIP, A. and RAMSAY, J. O. (2008). Combining registration and fitting for functional models. *J. Amer. Statist. Assoc.* **103** 1155–1165. MR2528838

KOENDERINK, J. J. (1990). *Solid Shape. MIT Press Series in Artificial Intelligence*. MIT Press, Cambridge, MA. MR1045203

KOENDERINK, J. and VAN DOORN, A. (1992). Surface shape and curvature scales. *Image Vis. Comput.* **10** 557–564.

MAO, Z., JU, X., SIEBERT, J. P., COCKSHOTT, W. P. and AYOUB, A. (2006). Constructing dense correspondences for the analysis of 3d facial morphology. *Pattern Recogn. Lett.* **27** 597–608.

MEIER, D. and FISHER, E. (2002). Parameter space warping: Shape-based correspondence between morphologically different objects. *IEEE Trans. Med. Imag.* **21** 31–47.

MEYER, M., BARR, A., LEE, H. and DESBRUN, M. (2002). Generalized barycentric coordinates on irregular polygons. *J. Graphics Tools* **7** 13–22.

OHTAKE, Y., BELYAEV, A. and SEIDEL, H.-P. (2004). Ridge-valley lines on meshes via implicit surface fitting. In *ACM Transactions on Graphics* (*TOG*) **23** 609–612. ACM.

PATRANGENARU, V. and ELLINGSON, L. (2016). *Nonparametric Statistics on Manifolds and Their Applications to Object Data Analysis*. CRC Press, Boca Raton, FL. MR3444169

PAULSEN, R. R. and HILGER, K. B. (2003). Shape modelling using Markov random field restoration of point correspondences. In *Biennial International Conference on Information Processing in Medical Imaging* 1–12. Springer, Berlin.

PENNEC, X., AYACHE, N. and THIRION, J.-P. (2000). Landmark-based registration using features identified through differential geometry. In *Handbook of Medical Imaging—Processing and Analysis* 499–513. Academic Press.

RAVIV, D., BAYRO-CORROCHANO, E. and RASKAR, R. (2017). LRA: Local rigid averaging of stretchable non-rigid shapes. *Int. J. Comput. Vis.* **124** 132–144. MR3681619

RAVIV, D. and KIMMEL, R. (2015). Affine invariant geometry for non-rigid shapes. *Int. J. Comput. Vis.* **111** 1–11. MR3300516

RAVIV, D., BRONSTEIN, A. M., BRONSTEIN, M. M., WAISMAN, D., SOCHEN, N. and KIMMEL, R. (2014). Equi-affine invariant geometry for shape analysis. *J. Math. Imaging Vision* **50** 144–163. MR3233140

ROHR, K. (2001). *Landmark-Based Image Analysis*: *Using Geometric and Intensity Models* **21**. Springer, Berlin.

RUSTAMOV, R. M., OVSJANIKOV, M., AZENCOT, O., BEN-CHEN, M., CHAZAL, F. and
    GUIBAS, L. (2013). Map-based exploration of intrinsic shape differences and variability. *ACM
    Trans. Graph.* **32** 72.

SRIVASTAVA, A., SAMIR, C., JOSHI, S. H. and DAOUDI, M. (2009). Elastic shape models for face
    analysis using curvilinear coordinates. *J. Math. Imaging Vision* **33** 253–265. MR2480989

STYLIANOU, G. and FARIN, G. (2004). Crest lines for surface segmentation and flattening. *IEEE
    Trans. Vis. Comput. Graph.* **5** 536–544.

VITTERT, L., BOWMAN, A. and KATINA, S. (2019a). Supplement A to "A hierarchical curve-based
    approach to the analysis of manifold data." DOI:10.1214/19-AOAS1267SUPPA.

VITTERT, L., BOWMAN, A. and KATINA, S. (2019b). Supplement B to "A hierarchical curve-based
    approach to the analysis of manifold data." DOI:10.1214/19-AOAS1267SUPPB.

VITTERT, L., BOWMAN, A. and KATINA, S. (2019c). Supplement C to "A hierarchical curve-based
    approach to the analysis of manifold data." DOI:10.1214/19-AOAS1267SUPPC.

VITTERT, L., BOWMAN, A. and KATINA, S. (2019d). Supplement D to "A hierarchical curve-based
    approach to the analysis of manifold data." DOI:10.1214/19-AOAS1267SUPPD.

VITTERT, L., BOWMAN, A. and KATINA, S. (2019e). Supplement E to "A hierarchical curve-based
    approach to the analysis of manifold data." DOI:10.1214/19-AOAS1267SUPPE.

L. VITTERT
A. W. BOWMAN
SCHOOL OF MATHEMATICS AND STATISTICS
UNIVERSITY OF GLASGOW
15 UNIVERSITY GARDENS
GLASGOW, G12 8QW
UNITED KINGDOM
E-MAIL: Liberty.Vittert@glasgow.ac.uk
          Adrian.Bowman@glasgow.ac.uk

S. KATINA
INSTITUTE OF MATHEMATICS AND STATISTICS
MASARYK UNIVERSITY
BRNO
CZECH REPUBLIC
E-MAIL: stanislav.katina@gmail.com