Curve Extraction and Facial Analysis Using Statistical Techniques

Jill Eileen Higgins

A Dissertation Submitted to the

University of Glasgow

for the degree of

Master of Science

Department of Statistics

December 2009

Abstract

The analysis of shape is an area of statistics which has wide applications, and is made particularly interesting by its visual nature. The field of facial shape analysis is particularly interesting because the human face has such a complex but familiar shape. This thesis investigates the facial characteristics of a cohort of 5-year-old children, whose faces were photographed, and the corresponding data captured, to provide a control group with which to undertake a study into Cleft Lip or Cleft Lip and Palate.

Chapter 1 provides an introduction to the data, in particular to the defined landmarks of the face, and to the shape analysis theory necessary for much of the analysis.

Chapter 2 involves the analysis of the landmark-based data, including consideration of the reliability with which landmarks are identified. The issue of sexual dimorphism is then introduced, and we consider whether any significant differences exist between the males and females of the cohort, including the use of principal components analysis to explore the data.

To expand upon the landmark-based analysis, the third and fourth chapters involve the development of methods with which to extract a set of facial curves.

Chapter 3 uses a combination of 2 methods for defining planes in three di-

mensional space with which to "cut" the face, and then identify the curve which lies at the intersection of the face and the plane. The set of facial curves identified using this method are then analysed using the same principal components analysis that was used for the landmark-based data.

Chapter 4 introduces theory and methodology to identify the facial curves based on the surface curvature characteristics of the face, and goes on to explore interesting features which were observed as part of the curve extraction.

Chapter 5 discusses the findings of the thesis and provides some thoughts on further work which would be of interest.

Acknowledgements

Firstly, I would like to thank my supervisor, Professor Adrian Bowman, for the guidance and support that he has provided during the year of my research and in the time since, particularly for the time that he has given up during the final period of my writing up.

I am also hugely grateful to the Statistics department of the University of Glasgow for the funding for my research, and for the tutoring and demonstrating opportunities I was afforded during my year in the department. I also feel very lucky to have been exposed to such varied and interesting areas of research.

I would like to thank Dr Sarah Barry for making the time to chat to me about her own research.

Thanks also go to my fellow postgraduate students, particularly those who I was lucky enough to share an office with, for helping me have a truly fun year and for putting up with my incessant chatter.

I would like to thank my family for being such a great support throughout my education, and for being a truly unique set of individuals!

Finally, thanks go to my husband Paul for always believing in me and making me smile.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Introduction

Shape analysis is an area of statistics with wide applications, and is made particularly interesting by its visual nature. The field of facial shape analysis is particularly interesting because the human face has such a complex but familiar shape. This thesis will investigate the facial characteristics of a cohort of 5-year-old children, whose faces were photographed, and data captured, to provide a control group with which to undertake a study into Cleft Lip or Cleft Lip and Palate.

This thesis can be considered in two sections. The first section, in Chapter 2, will consider the shapes of the faces of this cohort as defined by the landmarks of the face. The reliability of the landmark identification will be considered, and evidence of any sexual dimorphism will be sought.

The second section of the thesis, in Chapters 3 and 4, will consider two methods for extracting particular curves of the face, so as to provide an additional

starting point for analysis. Chapter 3 implements a method using planes in three dimensions to "cut" the face at appropriate angles, while Chapter 4 investigates the theory behind surface curvature in the context of the faces of the cohort, and used this theory to explore alternative methods of identifying one of the facial curves of interest.

1.2 Data

As part of facial analysis studies of children who had undergone surgery to repair a Unilateral Cleft Lip (UCL), or Unilateral Cleft Lip and Palate (UCLP), facial data was collected on a cohort of 100 control subjects, with no facial clefting, with the intention of capturing image data at the ages of 3 and 6 months, 1, 2 and 5 years. The initial analysis of the cohort, at 3 months of age, is discussed in White *et al.* (2004) The data to be considered in this study are the facial images of the subjects at 5 years of age, 88 of which were successfully captured. Using a computerized stereophotogrammetry (C3D) camera system, as discussed in Ayoub *et al.* (2003), three-dimensional images were obtained of the 88 children (42 males, 46 females), and 33 reproducible facial landmarks were identified for each subject by a single operator. The data available for analysis consists of around $19,000$ three-dimensional coordinates which define the surface of the face, although this value varies between subjects, from as low as $14,000$ to a maximum of $25,000$.

Also available are defined groupings of three coordinates, or 'triples', which allow the face to be represented as a triangular mesh, where each triple gives the coordinates of the vertices of a triangle within the mesh.

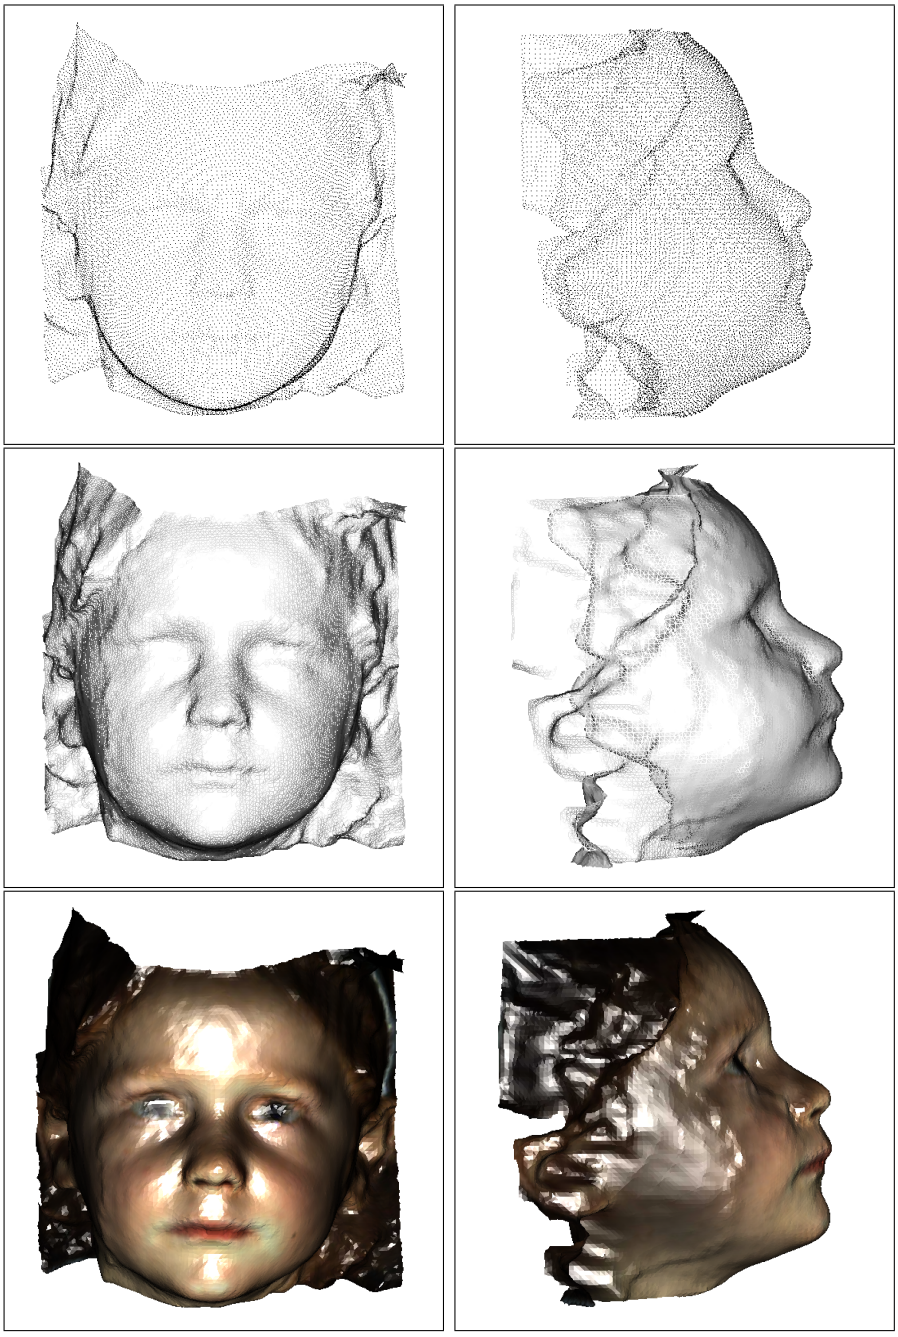

Figure 1.1: Facial data represented as three-dimensional coordinates (top),a three-dimensional mesh (centre) and a three-dimensional mesh with colour (bottom).

Further, for display and inspection purposes, the colour to be assigned to each point is available, and this in turn allows colour painting of each triangle. This allows the face to be viewed in a more easily recognisable form. Illustrations of the face in coordinate and mesh form can be seen in Figure 1.1.

1.3 Landmarks

A commonly used set of landmarks of the face are defined by Farkas (1994). Table 1.1 describes the subset of 33 soft-tissue landmarks which are used in this analysis, based mainly on Farkas (1994) with additional information from White *et al.* (2004). The landmarks are shown graphically in Figure 1.2. This subset includes most of the major landmarks of the face, however particular landmarks are omitted from this list, due to problems in accurate identification. An example of this is the *Otobasion Inferius*, which is defined as the point of attachment between the earlobe and the cheek, and is regularly not identified due to poor image quality towards the extremes of the facial area. The reliability with which landmarks are identified is examined in Section 2.1.1.

1.4 Areas for Investigation

Initial analysis will evaluate the degree of human error by using repeated landmark identification to identify any particular landmarks which are poorly extracted. The cohort will also be compared to a previous five-year-old non-cleft cohort to identify any irregularities between operators. Once the reliability of the data has been established, it is of interest to look for evidence of any sexual

Landmark Name	Landmark Abbreviation	Definition
Alar Crest	acL, acR	Lower limit point of the alar base
Alare	alL, alR	Most lateral point on the alar contour
Cheilion	chL, chR	Labial commissure point
Highest point of columella	cL, cR	Columella crest point level with top of nostril
Endocanthion	enL, enR	Inner eye fissure commissure point
Labiale inferius	li	Lower midline limit of vermillion
Labiale superius	ls	Upper midline limit of vermillion
Nasion	n	Midline point of both nasal root and nasofrontal suture
Pronasale	prn	Most prominent point of the apex nasi
Subalare	sbalL, sbalR	Lower limit point of alar base
Sublabiale	sl	Lower border of lower lip / Upper border of chin
Subnasale	sn	Midline junction point between columella and upper lip
Subnasale'	sn'L, sn'R	Midpoint of columella crest
Alare' inner	al'iL, al'iR	Midpoint of ala (inner)
Alare' outer	al'oL, al'oR	Midpoint of ala (outer)
Christa philtri	cphL, cphR	Junction point between upper lip vermillion and philtral peak
Exocanthion	exL, exR	Outer eye fissure commissure point
Highest point of nostril	hnL, hnR	Highest nostril point with head tilted at 60°
Stomion	sto	Midline contact point between upper and lower lip, or closest upper and lower lip points where lips incompetent
Upper lip border	ulbL, ulbR	Midpoint between ch and cph on the vermillion border

Table 1.1: Landmarks for Analysis

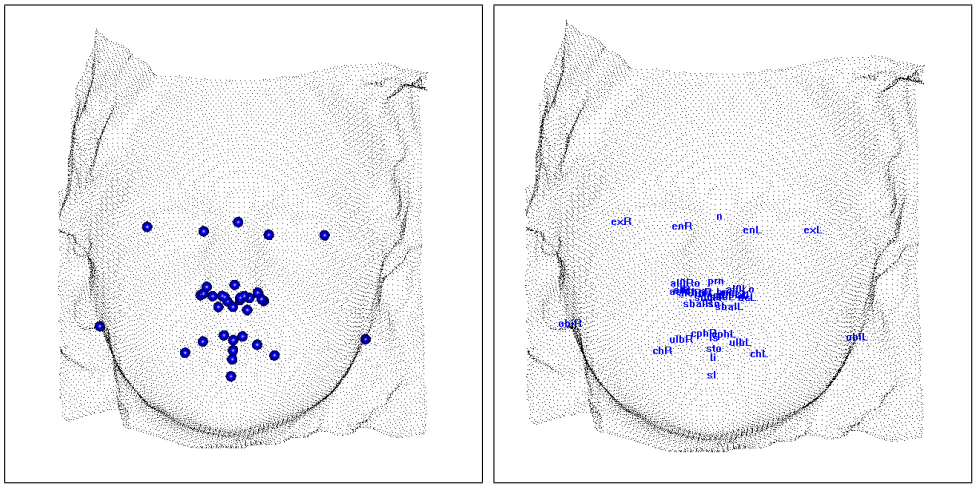

Figure 1.2: Facial landmarks identified on one subject.

dimorphism that may exist. Techniques will then be discussed for extracting particular facial curves, and it will then be possible to consider some more advanced analysis of the data.

1.5 Theoretical Considerations

The theory behind shape analysis based on landmark configurations is discussed in detail in Dryden & Mardia (1998) and the notation to be used follows from this. It is useful, firstly, to understand what is meant when discussing the concept of 'shape'. Dryden & Mardia (1998) define shape as "*all the geometrical information that remains when location, scale and rotational effects are filtered out from an object.*" Thus to compare the shape of two or more objects, it is appropriate to adjust the size, position in space and orientation of the objects so that they match as closely as possible to one another. Any discrepancies between these adjusted objects can therefore be interpreted as differences in shape.

The altering of the characteristics of an object in this way, and the subsequent analysis, is achieved through Procrustes Analysis.

1.5.1 Procrustes Analysis

Consider two $k \times m$ matrices, X_1 and X_2, each defining k landmarks in m dimensions and suppose that it is of interest to compare the shapes of the two objects defined by these configurations. Initially, it is useful to manipulate the data such that the columns of both X_1 and X_2 have been centred on 0. Full Ordinary Procrustes Analysis is used to rotate, scale and shift X_1 so that the matched configuration is as similar as possible to X_2. This is carried out by minimizing the squared Euclidean distance:

$$D^2{}_{OPA}(X_1, X_2) = \|X_2 - \beta X_1 \Gamma - \mathbf{1}_k \gamma^T\|^2$$

where $\|X\| = \sqrt{\text{trace}(X^T X)}$ (the Euclidean norm) and $\mathbf{1}_k$ is a vector of k $1s$. The minimization is over $\beta > 0$, which is a scale parameter; Γ, which is an $m \times m$ matrix describing the necessary rotation; and γ, an $m \times 1$ shift vector.
Using the matrix decomposition

$$X_2{}^T X_1 = \|X_1\| \|X_2\| V \Lambda U^T$$

where U and V are $m \times m$ rotation matrices and Λ is an $m \times m$ diagonal matrix, the value of $D^2{}_{OPA}(X_1, X_2)$ is then minimized by

$$\begin{aligned} \hat{\Gamma} &= UV^T \\ \hat{\beta} &= \frac{\text{trace}(X_2{}^T X_1 \hat{\Gamma})}{\text{trace}(X_1{}^T X_1)} \\ \text{and } \hat{\gamma} &= \mathbf{0}_m \text{ (a vector of } m \text{ zeros)} \end{aligned}$$

The full Procrustes fit of X_1 onto X_2 is then defined as

$$X_1{}^P = \hat{\beta} X_1 \hat{\Gamma} + \mathbf{1}_k \hat{\gamma}^T.$$

A partial implementation of this idea can be used to compare two landmark configurations when manipulation of the size of the objects is unnecessary. Further, a generalized version of Procrustes Analysis can be used when $n \geq 2$ configuration matrices are available, and this is outlined in Section 1.5.3.

1.5.2 Evaluating Asymmetry

To ascertain the degree of asymmetry of a shape, a reflected and relabeled configuration can be compared to the original, and the extent of the mismatch provides a measure of asymmetry. The configuration for analysis can be represented by the $k \times m$ matrix, X, defining k landmarks in m dimensions. A further, reflected landmark configuration can then be described by the $k \times m$ matrix, X_R, which is constructed by reversing the first-dimension coordinate for each landmark and then by reversing the labeling of each pair of corresponding landmarks. Thus 'left'-labeled landmarks are relabeled as 'right', and vice versa, while centrally positioned landmark names remain unchanged. The positioning of the shapes defined by X and X_R in space are not of interest in quantifying asymmetry, and so it is assumed that the columns of X and X_R have been centred on 0. Further, since the two matrices for comparison are reflections of one another, we can assume for convenience that $\|X\| = \|X_R\| = 1$.

The configurations X and X_R are then matched as closely as possible to one another using Ordinary Partial Procrustes Analysis to minimise the squared Eu-

clidean distance

$$D^2{}_{OPA}(X, X_R) = \|X - X_R\Gamma - \mathbf{1}_k\gamma^T\|^2$$

where Γ is an $m \times m$ rotation matrix, γ is an $m \times 1$ shift vector and $\mathbf{1}_k$ is a vector of k $1s$. Note that the scale parameter β is set to 1, as an alteration of scale is unnecessary since the two configurations have unit size.

Using the matrix decomposition

$$X^T X_R = \|X\|\|X_R\|V\Lambda U^T$$

where U and V are $m \times m$ rotation matrices and Λ is an $m \times m$ diagonal matrix, the value of $D^2{}_{OPA}(X, X_R)$ is then minimized by

$$\begin{aligned} \hat{\Gamma} &= UV^T \\ \text{and } \hat{\gamma} &= \mathbf{0}_m \text{ (a vector of } m \text{ zeros)} \end{aligned}$$

Thus to match the configurations as closely as possible, all that is necessary is the rotation defined by $\hat{\Gamma}$. No shifting or scaling is necessary in the case of comparing reflected, centred images.

A measure of asymmetry, as proposed by Bock & Bowman (2006), can then be defined by

$$A = \frac{\|X - X_R\hat{\Gamma}\|^2}{k}$$

which quantifies the degree of mismatch between X and X_R, standardized by the number of landmarks so as to represent the average displacement per landmark.

1.5.3 Comparing the Mean Shape of Two Populations

It is desirable to be able to estimate the mean shape of a population, and further to be able to compare two population mean shapes, and indeed to test for

equality. All theory and notation is again taken from Dryden & Mardia (1998).

Suppose that a sample of n subjects is available. The $k \times m$ configuration matrix for subject i, where $i = 1, \ldots, n$ is denoted by X_i, giving k defined landmarks in m dimensions. Within a population, to estimate the mean shape, it is necessary to remove the effects of location, scale and orientation. Initially, it is again useful to manipulate the data such that the columns of each X_i have been centred on 0, and that the scale of each configuration is set to 1.

To compare $X_1, \ldots, X_n$, Generalized Procrustes Analysis is used to minimize, over the β_i, Γ_i and γ_i,

$$\frac{1}{n} \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} \|(\beta_i X_i \Gamma_i + \mathbf{1}_k \gamma_i^T) - (\beta_j X_j \Gamma_j + \mathbf{1}_k \gamma_j^T)\|^2 \tag{1.1}$$

subject to the constraint that

$$\bar{X} = \frac{1}{n} \sum_{i=1}^{n} (\beta_i X_i \Gamma_i + \mathbf{1}_k \gamma_i^T)$$

has size 1. The minimum value of (1.1) is then denoted by $G(X_1, \ldots, X_n)$, the Generalized Procrustes Sum of Squares, which is found by a numerical algorithm which defines the $\hat{\beta}_i$, $\hat{\Gamma}_i$ and $\hat{\gamma}_i$, the minimizing parameters.

The resulting fitted configurations are then defined by ${X_i}^P = \hat{\beta}_i X_i \hat{\Gamma}_i + \mathbf{1}_k \hat{\gamma_i}^T$ for $i = 1, \ldots, n$ in the shape space.

Once the X_i have been fully Procrustes matched, the arithmetic mean shape is obtained by simply taking the mean of each coordinate for each landmark, so that $\bar{X}^P = \frac{1}{n} \sum_{i=1}^{n} {X_i}^P$ is the full Procrustes mean. The tangent space to a shape space is a concept which provides a linear approximation to the shape space so that multivariate techniques can be implemented. In fact, the Procrustes residuals ${X_i}^P - \bar{X}^P$ provide a good approximation to the tangent space coordinates, provided the shapes in the sample are fairly similar.

To compare two independent populations with mean shapes $\bar{X}_1^P$ and $\bar{X}_2^P$ estimating population means μ_1 and μ_2, a Hotelling's T^2 two sample test can be carried out for

$$H_0 : \mu_1 = \mu_2 \text{ against } H_1 : \mu_1 \neq \mu_2$$

in the Procrustes tangent space calculated from the combined populations.

Chapter 2

Preliminary Analysis of Landmark Data

The first stage in analysis of the facial data is to explore any sources of error in the data collection. Ayoub *et al.* (2003) discusses sources of error due to the nature of the camera system. Small capture errors occur due to an instability in the C3D system when multiple images are taken, and discrepancies also occur when differences arise in where the object for image capture is placed in relation to the cameras. These errors are comparable with other 3D image capture systems (Ayoub *et al.* (2003)). Marking up of landmarks is a source of error which can be quantified for this particular set of data, firstly by considering the degree to which the operator was able to reproduce the same landmark coordinates (see Section 2.1.1), and secondly by comparing the mean landmark configuration shape of this cohort of control 5-year-old subjects with a previous 5-year-old cohort (see Section 2.1.2.

2.1 Reliability in Landmark Identification

2.1.1 Single Operator Consistency

In order to assess the degree of reliability with which each landmark was identified, for 12 subjects the facial landmarks were identified on three separate occasions, and these repeated identifications were examined to establish consistency. Since the comparisons were for the same images each time, it was not necessary to use Procrustes Analysis. The mean coordinate value for each landmark, for each subject, was calculated from the three repeated identifications, and the reliability of the repeated identifications was quantified by the Euclidean distance between each identified landmark and the corresponding mean value. Each landmark was considered individually, so that any landmark which was significantly inconsistently identified could be determined. For each of the 12 subjects, there were 3 measures of deviation from the mean value for each of the 33 landmarks, so for each landmark 36 ($= 3 \times 12$) values were available for assessment of reliability in identification. The results are shown in Figure 2.1.

The mean operator error was found to be 0.405mm, across all landmarks, and 96.5% of the landmarks were identified as being within 1mm of the mean, with only 3 of the 1180 ($= 33 \times 36$) distances being greater than 1.5mm. These landmark identification errors are greater than the errors observed by Ayoub *et al.* (2003) on a cohort of infants with cleft lip at the time of the first surgical treatment (which occurs between 3 and 6 months), however the values are still small, and so it seems acceptable to assume that the errors generated by human inconsistency in landmark identification are sufficiently small that data analysis should proceed. It is also worth noting at this point that the previous analysis was

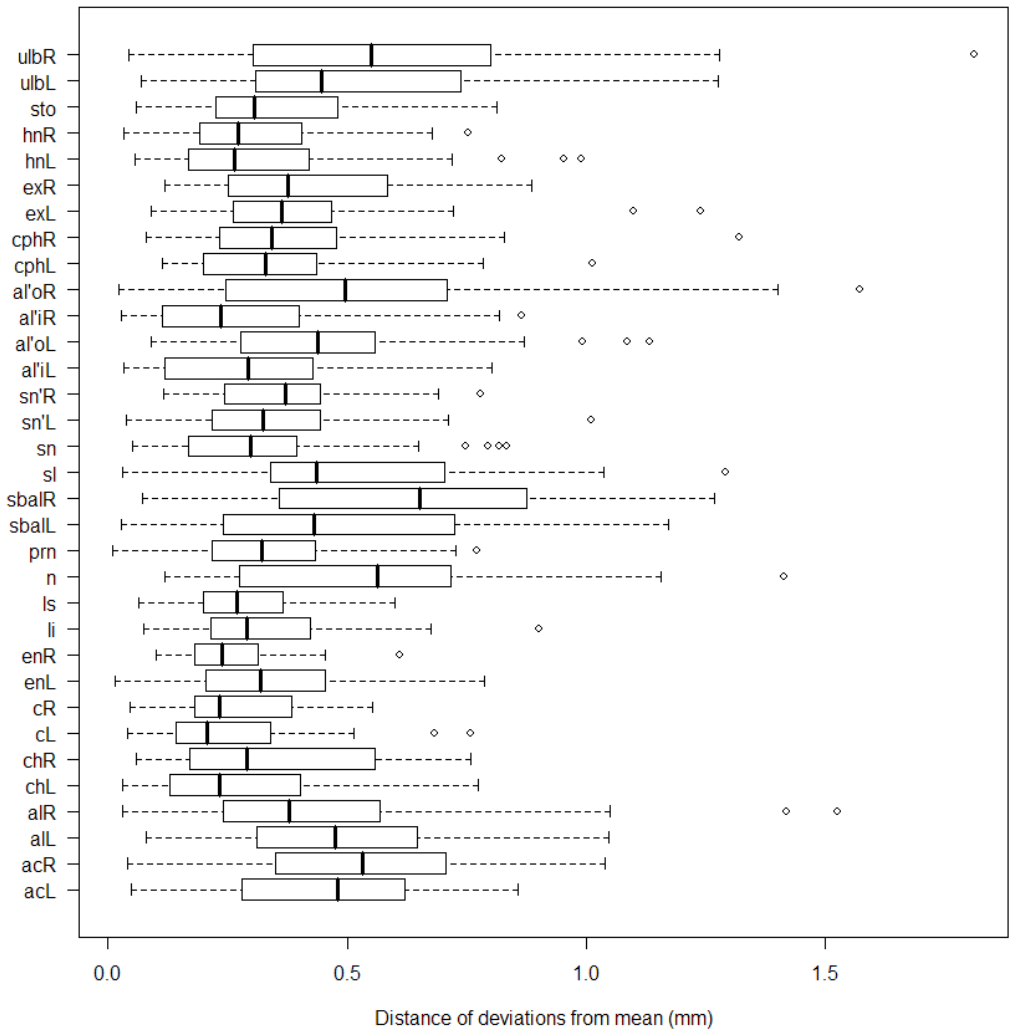

Figure 2.1: Landmark Identification Consistency

on much younger infants, with correspondingly smaller faces, and so we would perhaps expect that absolute measurement errors would be larger.

2.1.2 Consistency between Different Operators

The sets of identified landmarks were compared with previous data from a different cohort of five-year-old control subjects, which had been identified by a different single operator. Generalized Procrustes Analysis (as discussed in Section 1.5.3) was carried out on a subset of the landmarks to test the two populations (cohort of interest and previous cohort) for equality of mean shape. An overall difference in mean shape was noted, which was an unexpected finding, so it was of interest to establish where these differences lay by carrying out a Multivariate Analysis of Variance (MANOVA) for each individual landmark within the Procrustes matched configurations. The findings are shown in Table 2.1. It can be seen that the majority of the Procrustes-matched landmarks are found to be from significantly different populations, although the differences are small in size, and so do not provide a major cause for concern. However, there is no obvious reason as to why this should be the case, and so it is interesting to consider why the two cohorts should differ. One explanation could be that there are consistent differences between the operators in the interpretation of landmark definitions, which would lead to consistent differences in the marked-up landmark coordinates. Alternatively, the two sets of data may have been sampled from populations with slightly different characteristics, although it is not obvious why these differences would occur. We should also exercise caution, however, in interpreting the MANOVA results since the sample size is relatively small, and indeed is not much larger than the number of dimensions.

Landmark	Significant Cohort Difference	Landmark	Significant Cohort Difference
n	*	chL	
acR	*	chR	
acL	*	ls	
enR		sl	
enL	*	li	*
prn	*	exL	
alL	*	exR	
alR	*	cL	*
sn	*	cR	
sbalL	*	hnR	*
sbalR		hnL	*
ulbL	*	cphL	*
ulbR	*	cphR	
al'iL	*	al'oL	*
al'iR	*	al'oR	*
sn'L	*	sto	*
sn'R	*		

* Significant Difference shown at 1% level

Table 2.1: Comparing Historical and Current Operator Landmark Identification

It would be interesting to be able to consider whether any difference in populations were noted under either of the following circumstances:

- The landmarks of the two cohorts were to be marked up by the same operator
- Two different operators were to mark up a single cohort

Either of these analyses may provide further insight into why population differences were noted in this case.

A two-dimensional frontal snapshot of the three-dimensional superimposed Procrustes matched mean shapes is shown in Figure 2.2. Note that the mean shapes are very similar and so if only one point is visible at a particular landmark, this implies that the corresponding mean landmark from the other cohort is hidden behind the visible mean landmark at this particular angle. A frontal view of the landmark configurations of all subjects after full Procrustes matching can be seen in Figure 2.3.

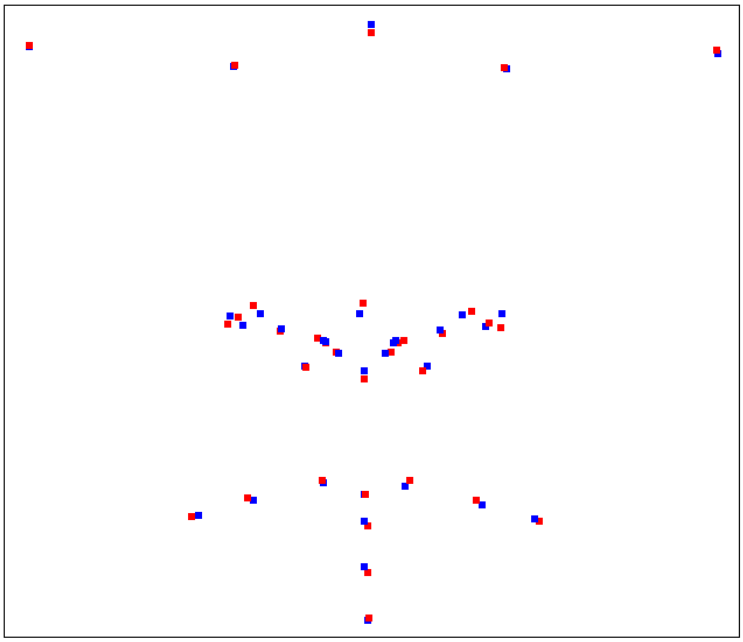

Figure 2.2: Frontal view of the full Procrustes mean landmark configurations of the current cohort (Red) and previous cohort (Blue), rotated to each other using GPA.

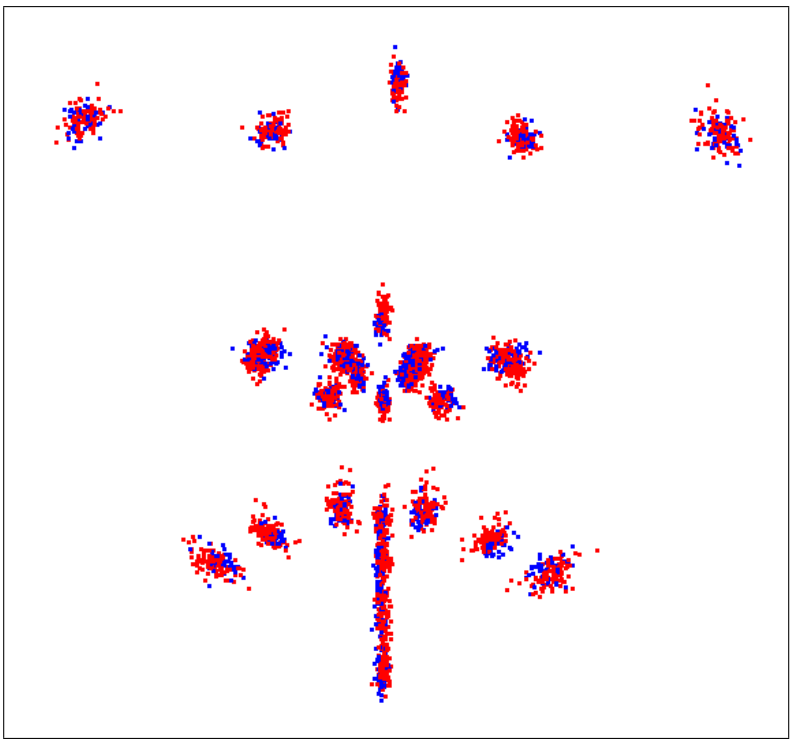

Figure 2.3: Frontal view of the full Procrustes individual subject landmark configurations of the current cohort (Red) and previous cohort (Blue), rotated to each other using GPA.

2.2 Sexual Dimorphism

It is of interest to consider whether any statistically significant differences exist between males and females of the current cohort. Previous studies have indicated that certain inter-gender size-based facial differences do exist at 3 months (White *et al.* (2004)), 4 months to 3.5 years (Yamada *et al.* (2002)), 8 to 14 years (Ferrario & Sforza (1997)) and in adults (Ferrario *et al.* (1995)). White *et al.* (2004) found that differences in facial measurements could be explained by differences in the weight of the subjects, and Ferrario *et al.* (1995) and Ferrario & Sforza (1997) both noted that once the effect of size was eliminated no significant facial sexual shape dimorphism existed. Yamada *et al.* (2002), however, found no evidence of a correlation between the size of the face and body size within any of the three age groups studied.

2.2.1 Basic analysis

It is differences in shape that are the main focus of interest here, but initially, it was of interest to obtain a measurement of overall size to assess any difference. This was obtained by calculating the sum of the distances from each landmark to the centroid of the configuration, and then scaling by the number of landmarks so as to provide an indication of the mean distance from landmark to centroid. The results of the comparison between male and female size (scaled by a factor of 100 for ease of interpretation) are shown in Figure 2.4. The boxplots represent the mean $\pm$ twice the standard deviations of the respective groups. It is clear to see that the mean sizes of the male and female configurations are very close, but that the spread of configuration sizes in males extends to noticeably higher values. A

two sample Student's t test found this difference in size to be significant at the 5% level, which is in agreement with previous literature. Once this effect of size is removed, however, it is of interest to examine the shape of the configurations and test for any sexual dimorphism. This was carried out by Generalized Procrustes

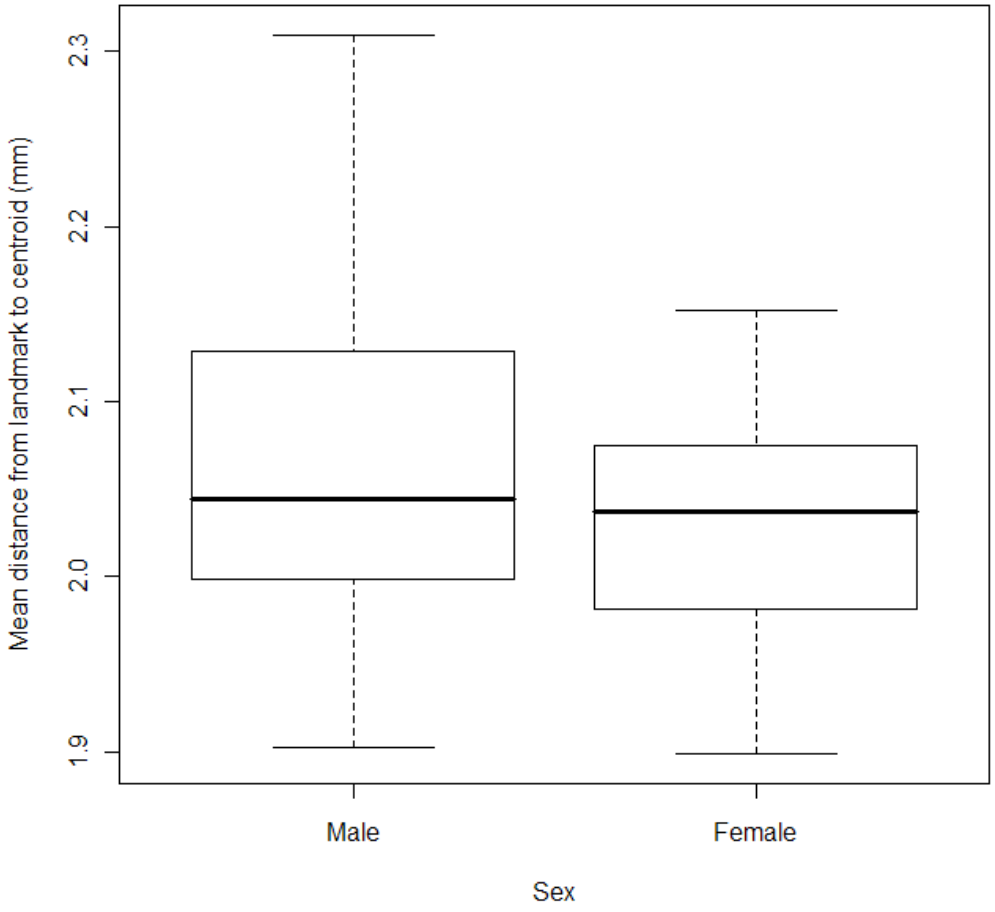

Figure 2.4: Comparison of Male and Female Size

Analysis and the Hotelling's T^2 two sample test as described in Section 1.5.3. No significant shape difference was found to exist between the males and females of the study. A two-dimensional frontal snapshot of the three-dimensional superimposed Procrustes matched mean shapes is shown in Figure 2.5. As in the previous section, where only one (male or female) mean landmark is visible, the corresponding (female or male) mean landmark is hidden at this particular angle.

A frontal view of the landmark configurations of all subjects after full Procrustes matching can be seen in Figure 2.6.

2.2.2 Principal components analysis

An additional interesting way in which the male and female landmark configurations can be compared is through principal components analysis. This is a type of multivariate analysis which allows the reduction of the dimensionality of the data. As it stands, the facial data has a high number of landmarks, each defined in three dimensions. Principal components analysis seeks to identify a small number of uncorrelated linear combinations of the variables which explain most of the variation in the data. The theory and notation to follow are taken from Venables & Ripley (1997).

Consider an $n \times p$ matrix of data, X, and let S denote the covariance matrix of X. Then S is defined as follows:

$$nS = (X - n^{-1}\mathbf{1}\mathbf{1}^T X)^T (X - n^{-1}\mathbf{1}\mathbf{1}^T X) = (X^T X - n\bar{\mathbf{x}}\bar{\mathbf{x}}^T)$$

where $\bar{\mathbf{x}} = \mathbf{1}^T X/n$ is the row vector of the means of each of the variables.

To obtain the desired linear combinations with maximal (or minimal) variance, we can consider the sample variance of a linear combination $\mathbf{xa}$ of a row

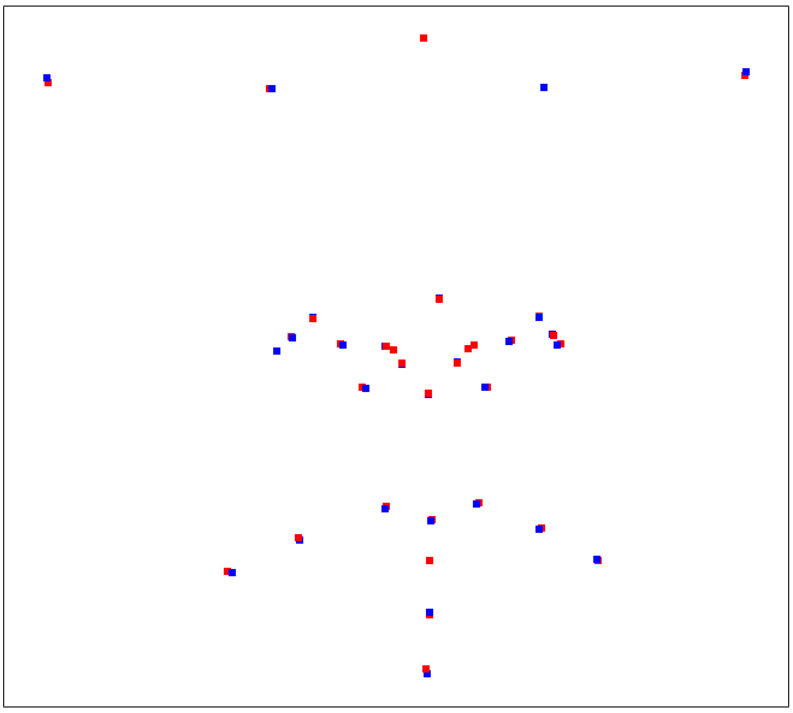

Figure 2.5: Frontal view of the full Procrustes mean landmark configurations of the male (Red) and female (Blue) subjects, matched using GPA.

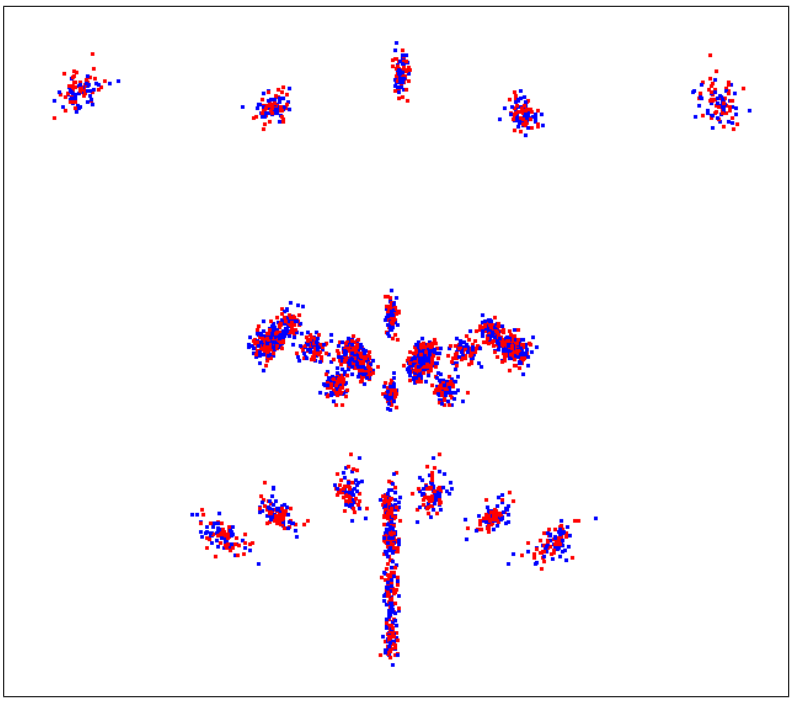

Figure 2.6: Frontal view of the full Procrustes individual subject landmark configurations of the male (Red) and female (Blue) subjects, rotated to each other using GPA.

vector $\mathbf{x}$, which is $\mathbf{a}^T\Sigma\mathbf{a}$. This is then maximised (or minimised) subject to $\|\mathbf{a}\|^2 = \mathbf{a}^T\mathbf{a} = 1$, since scaling of the variance is straightforward by rescaling the linear combinations.

Since Σ is a non-negative definite matrix, it has an eigendecomposition

$$\Sigma = C^T\Lambda C$$

where Λ is a diagonal matrix of (non-negative) eigenvalues, which are in decreasing order. Let $\mathbf{b} = C\mathbf{a}$, then we require to maximise $\mathbf{b}^T\Lambda\mathbf{b}(= \sum \lambda_i b_i^2)$ subject to $\sum b_i^2 = 1$. We can take $\mathbf{a}$ to be the column eigenvector corresponding to the largest eigenvalue of Σ, and subsequent eigenvectors give combinations with as large as possible variance but which are uncorrelated with those that have been taken earlier. The i^{th} principal component is then the i^{th} linear combination obtained using this method.

We can implement this theory to look at the principal components of the Procrustes-matched facial landmarks data. Figure 2.7 shows the cumulative proportion of the variability in the data which is explained by successive principal components. It can be seen that each successive principal component explains less of the variability than the previous principal component, as we would expect from the definition. The first 20 principal components are necessary to explain 90% of the variability. The first two together explain 39% of the variability.

It is of interest to look at the relationship between the (uncorrelated) first and second principal component scores, which are realisations of the principal components, and to consider whether any significant difference exists between the scores of the males and females of the cohort. Figure 2.8 shows boxplots of the first and second principal component scores, split by sex. Although the

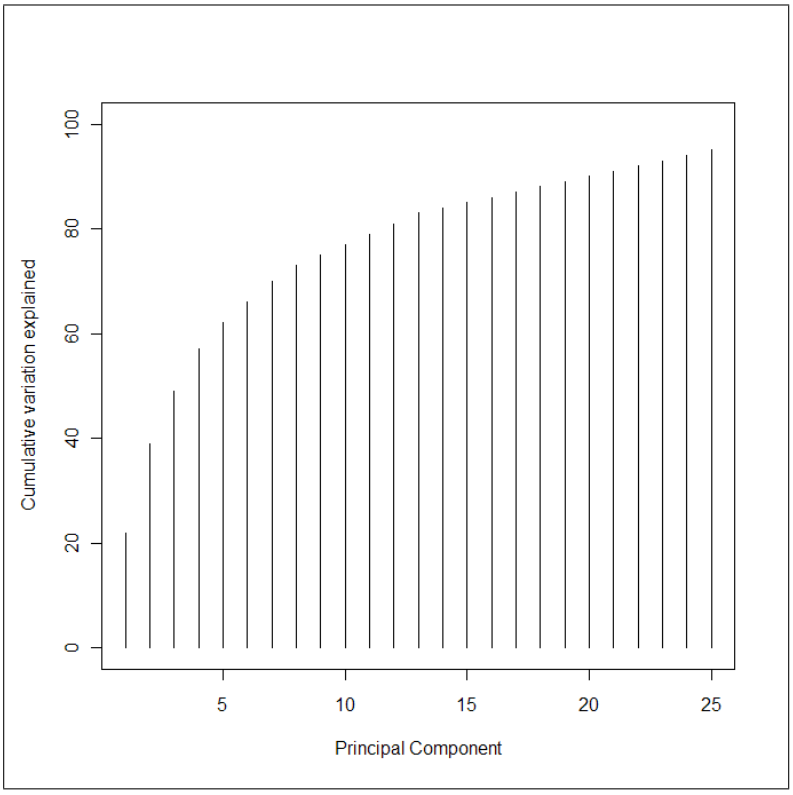

Figure 2.7: Cumulative proportion of the variability in the data explained by each additional principal component.

spreads of the scores look visually different for the first principal component score, a Student's t-test found no significant difference at the 5% level between the principal component scores of males and females for the first or second principal component scores. Indeed, the vast majority of later principal component scores are also not significantly different between males and females.

Figure 2.9 plots the first and second principal component scores against one another, split by males and females. As would be expected from the previous analyses, there is no obvious split between males and females.

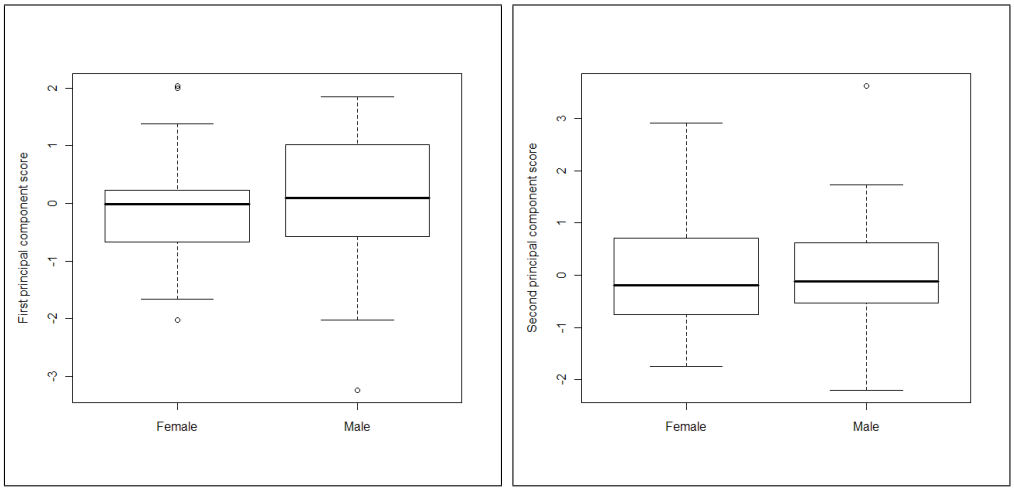

Figure 2.8: Boxplots of the first and second principal component scores, split by sex.

The findings from both sexual dimorphism analyses are in agreement with the findings of White *et al.* (2004) that the variation between the sexes at these early ages is based on size, rather than shape. It should be noted, however, that other studies have observed shape differences between the sexes, such as Ferrario & Sforza (1997) on 8 to 14 year-old subjects and Prahl-Andersen *et al.* (1995) on 9 to 14 year-old subjects. However, these studies are based on older children. It should also be noted that the analysis in this case is based on landmarks. When the analysis is extended to include the curvature of the face, the results may differ.

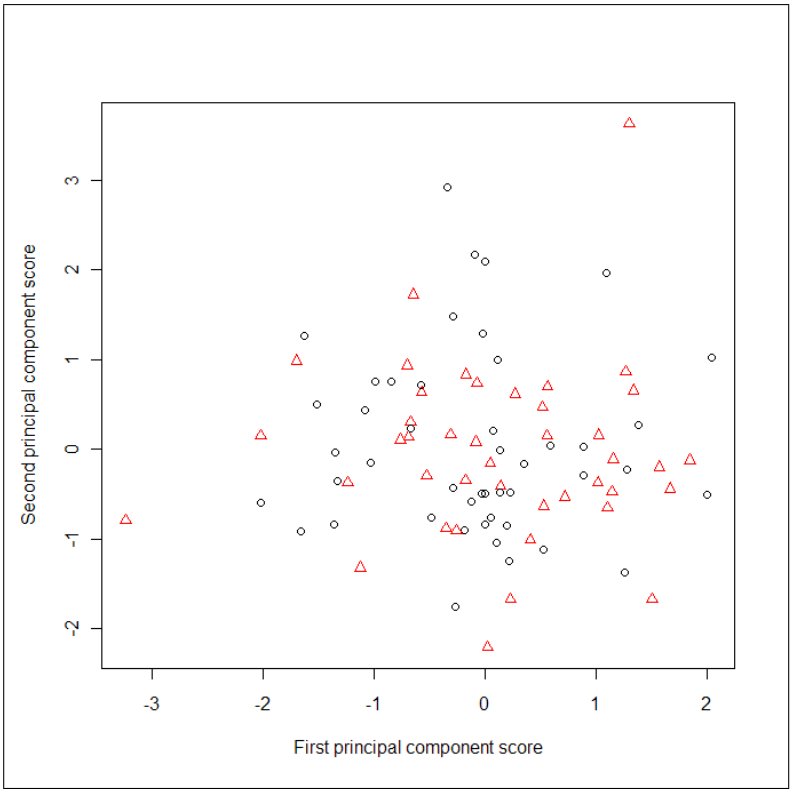

Figure 2.9: Plot of the first against second principal component scores, comparing males (triangles) against females (circles).

Chapter 3

A Planar Method for Facial Curve Identification

Chapter 2 provides a preliminary analysis of the data, based on coordinates of pre-defined landmarks. This yields useful results; however, since it uses only the landmarks, a substantial proportion of the available facial information is not being used. To develop the analysis, it is therefore desirable to extract further information from the facial data. Previous studies, such as Bowman & Bock (2006) and Bock & Bowman (2006) have used facial curves to examine the data. The aim of this chapter and the next is to explore methods with which curves of the face can be identified. This thesis will not attempt to analyse these curves, but will concentrate on methods for their identification. In this context, the approach adopted in identifying a facial curve is to identify a series of points which are close enough together so that they effectively define a curve. This chapter will discuss a planar method of curve extraction using the defined landmarks, and Chapter 4 will then discuss a method of curve extraction which uses the curvature of the

facial surface.

The curves to be extracted in this chapter are listed below. In some cases, curves are split into two or more sections for extraction, as the complexity of the curve prevents accurate capture in a single plane. In particular, the relatively complex shape of the upper lip curve involves splitting this into 6 sections for curve extraction.

Figure 1.2 provides a graphical illustration of the locations of these landmarks.

- Eye Curve
    1. “enR” to “n” to “enL”
- Midline Curve
    1. “n” to “prn” to “sn”
    2. “sn” to “ls”
- Nasal Bridge
    1. “acR” to “al’oR”
    2. “al’oR” to “prn” to “al’oL”
    3. “al’oL” to “acL”
- Nasal Base
    1. “sbalR” to “sn” to “sbalL”
- Lip Curve
    1. “chR” to “ulbR”

2. “ulbR” to “cphR”
3. “cphR” to “ls”
4. “ls” to “cphL”
5. “cphL” to “ulbL”
6. “ulbL” to “chL”

The first stage in extracting a facial curve is to define the plane in which the curve is to lie. Two complementary methods for defining the plane have been utilized, and these are discussed in Section 3.1. For illustration purposes, consider the Midline Curve, which runs from the top of the nose to the midpoint of the upper lip; from nasion (“n”), through pronasale (“prn”) and subnasale (“sn”) to the labiale superius (“ls”). In the case of the Midline Curve, the required plane would therefore be that which effectively bisects the face. Three views of this plane can be seen in Figure 3.1

Using simple algebra, it is then straightforward to identify which of the facial mesh line segments cross the plane, and to extract only those lying between the relevant landmarks. In the case of the Midline Curve, this will give a series of line segments, all of which were originally part of the overall facial mesh, and crossing the facial bisecting plane. Again using only elementary three-dimensional geometry, for each chosen segment, the particular point on the line which intersects with the plane can be identified. In the case of the Midline Curve, this will result in a series of irregularly spaced points, running down the ridge of the nose from the *nasion* to the *labiale superius*. For analysis of these curves, it is desirable to regularize the points so that the curve information for each subject contains the same number of regularly spaced points. The final stage in curve extraction is

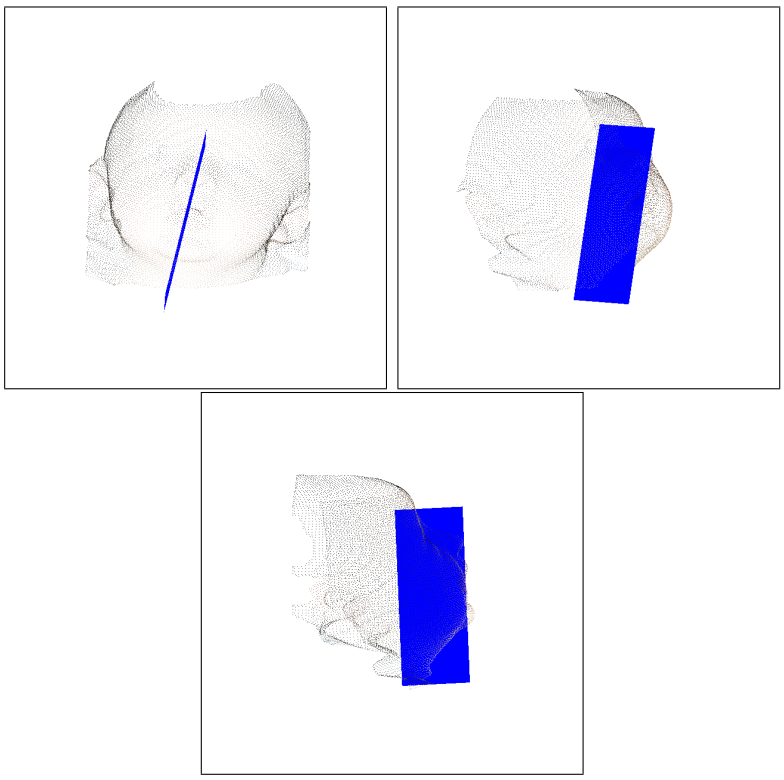

Figure 3.1: The plane which passes through the midline of the face.

therefore to fit a cubic smoothing spline to the extracted points. The points are defined in terms of the distance along the curve, t_i, in the form $(x(t_i), y(t_i), z(t_i))$, and smoothing is therefore carried out for $x(t_i)$ against t_i, $y(t_i)$ against t_i and $z(t_i)$ against t_i separately. For each smoothing step, this essentially involves finding the function $f(x)$, with two continuous derivatives, to minimize the penalized sum of squares

$$\sum_{i=1}^{n}\{y_i - f(x_i)\}^2 + \lambda \int_a^b \{f''(t)\}^2 \mathrm{d}t,$$

where λ is a constant and $a \leq x_1 \leq \ldots \leq x_n \leq b$. The effect of this minimization is that the function is chosen by its closeness to the data, but with a penalty

applied for curvature, in order to select a sufficiently smooth function. Predicted values from this function are then identified at regularly spaced intervals to generate regularly spaced points for analysis. The theory behind fitting the cubic smoothing spline is taken from Hastie & Tibshirani (1996). Figure 3.2 shows the process of curve extraction for the first part of the Midline Curve.

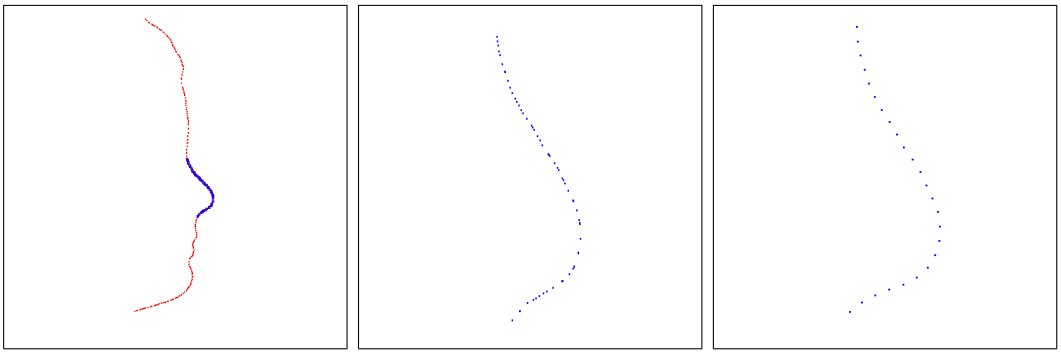

(a) (b) (c)

Figure 3.2: (a) Points of intersection between the triangular facial mesh and the plane which passes through the nasion, pronasale and subnasale, with required points shown in blue, (b) Close-up view of the irregularly spaced points extracted and (c) regularized points appropriate for analysis.

3.1 Defining the Plane

Method 1 - Three Planar Landmarks

There exists a unique plane which passes through any three given points, so for example, the plane which is to include the Midline Curve can be defined as that which passes through "n", "prn" and "sn". Unfortunately, it emerges that Method 1 is not appropriate for defining all desired facial curves. An example of

a curve which is not extracted well using Method 1 is the curve running across the upper lip, from the right to left *cheilion* ("chL" to "chR"). It appears that the reason for this is that the plane which passes through, for example, the set of three points ("cphL", "ulbL", "chL") is actually at an angle very close to parallel to the face. Hence the plane cuts not only the required facial curve, but also other facial curves as well. In the particular case of the upper lip, a curve very close to the required one is also cut (running through the surface within the upper lip), and this leads to problems in tracking only the required points, as can occur in the first image in Figure 3.3. It seems necessary to define a plane which will cut the face at a more appropriate angle.

Method 2 - Two Planar Landmarks and a Reference Landmark

As an alternative to using the plane defined by three landmarks, as in Method 1, we can define the plane by using the two landmarks which are to be the endpoints of the curve as lying in the plane, under the constraint that this plane must have a third "reference" landmark orthogonal to the plane. This allows the definition of a plane which will intersect the face in a manner which ensures that it is closer to perpendicular to the face than parallel. However, Method 2 can only be implemented if there exists a landmark which can provide a reliable orthogonal direction. If there does not exist such a landmark, then a "closest-match" is chosen, which is a landmark as closely as possible orthogonal to the required plane. However, this can lead to the tracking of a curve with the correct endpoints, but not following the required path. An example of when Method 2 does not give the correct curve is with the first part of the Midline Curve. The closest landmark found to being orthogonal to the desired plane, was found, by repeated examination, to be the alare ("acL"), but the result of using Method 2 is

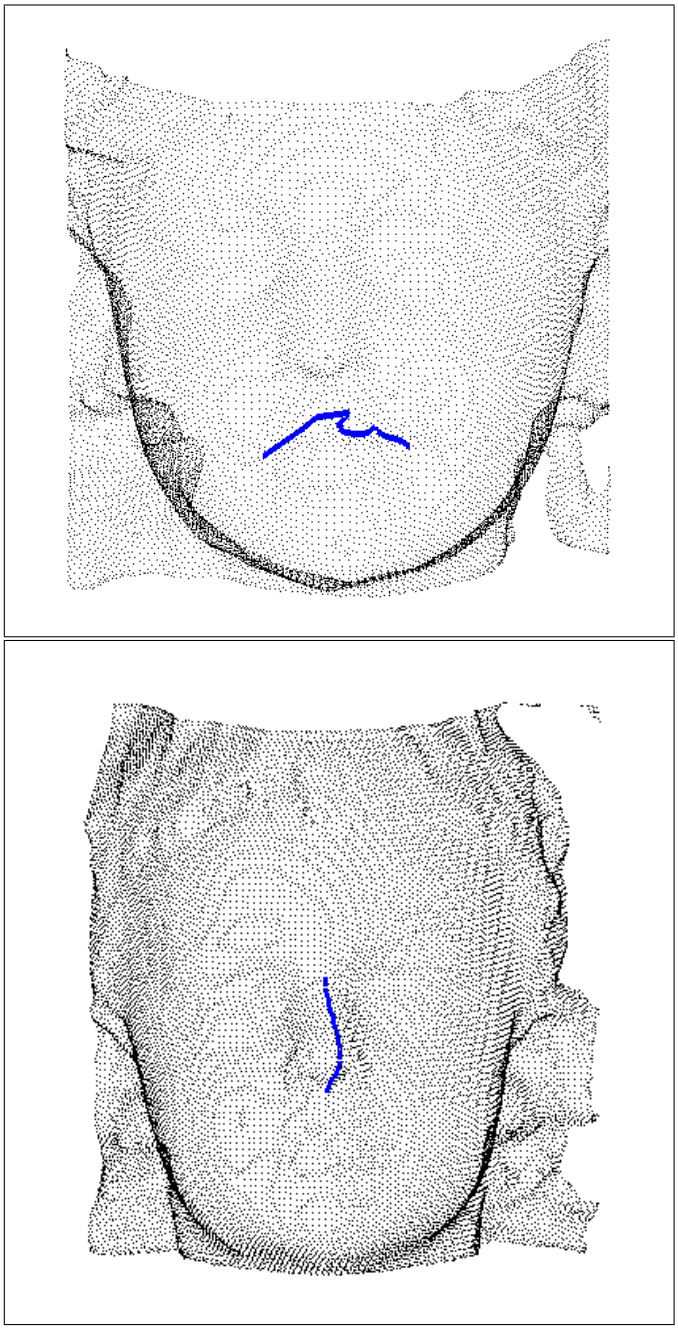

Figure 3.3: Failure of method 1 to correctly extract the upper lip curve and failure of method 2 to correctly extract the upper midline curve.

the second curve in Figure 3.3, which clearly does not run through the pronasale landmark.

3.2 Discussion on Implementation of Methods 1 and 2

To extract the facial curves, the first step is to choose which of Method 1 or 2 should be used to define the plane. The choice of method which we will implement is based on comparisons which have been made on the relative performances of each method for different parts of the face. The algorithm has therefore been set up to use the method which has been found to be most effective for each different curve (or section of curve). Once the plane has been defined, the same method for curve extraction is used, regardless of the planar definition method. It is unfortunate that it has not been possible to effectively use only one of Method 1 or 2 to extract the full set of facial curves. However, a combination of the two methods is effective in extracting a full set of facial curves, as seen in Figure 3.4.

Table 3.1 details which method was used for the extraction of each of the facial curves.

The use of planes to specify a two-dimensional curve describing a particular facial characteristic can be seen to be effective, and can be used to identify each of the facial curves of interest in the context of the cleft lip and/or palate studies for which the data was collected.

Limitations of the method are acknowledged, in particular:

- Landmarks are identified by a trained operator, however potential irregularities arise (as discussed in Section 2.1), which can affect not only the

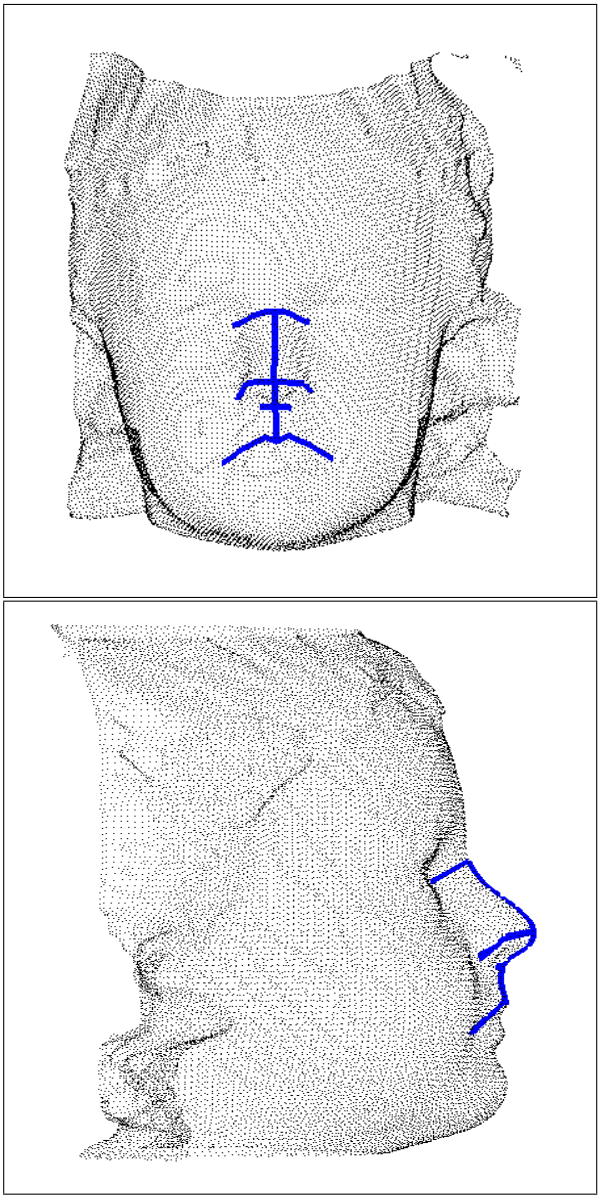

Figure 3.4: Full set of extracted curves using a combination of methods 1 and 2.

Curve	Method for defining the plane
Midline	Method 1 (three planar landmarks)
Eye curve	Method 1 (three planar landmarks)
Nose	Method 2 (two planar landmarks, one orthogonal landmark)
Base of nose	Method 1 (three planar landmarks)
Lip	Method 2 (two planar landmarks, one orthogonal landmark)

Table 3.1: Methods for defining the plane for each facial curve

endpoints of the curve, but any additional reference points defining the plane on which the curve is constrained to lie.

- There is no good reason why each facial curve should lie on a single plane. In particular, in the context of facial analysis, it is of interest to identify any unusual features, which may have underlying curves which deviate strongly from a single plane.

- It does not seem efficient to split curves into multiple sections for identification, and although it has not been a problem in the control cases for which this method has been tested, this method could result in curves which have obviously discontinuous derivatives at the point of intersection of the sections, particularly when unusual facial features are considered.

- Often the curves generated using this planar method are simply a series of points connecting landmarks to one another in the most direct way across the surface. It would be of great interest to identify those curves which are

inherent to the face, for example the 'ridge curves' of the face, which are discussed in Section 4.2

However, as can be seen in the illustrative curves in Figure 3.4, the method of using planes to "cut" the face and hence identify curves does produce a set of plausible curves, and indeed, the planar constraint imposed on the curves may make any analysis or comparison more straightforward.

Chapter 4 explores alternative strategies, based on the surface curvature characteristics of the face, for defining the facial curves.

3.3 Principal components analysis on the extracted curves

In Chapter 2, we concluded by carrying out principal components analysis on the landmark data to look for the existence of any significant differences between the males and females of the cohort. Now that we have full sets of extracted facial curves for each individual, the natural progression is to repeat this analysis but with the curve data, which is of higher dimension, and correspondingly includes richer detail about the faces. The high dimensionality also means that the data is particularly appropriate for principal components analysis, which aims to reduce dimensionality. The analysis of the curves was carried out in the same way as the landmark analysis, with the smoothed points that characterise the curve being used as pseudo-landmarks.

Figure 3.5 shows the cumulative proportion of the variability of the data explained by successive principal components. The first 15 principal components are necessary to explain 90% of the variability, compared to the first 20 for the

landmark data. Together, the first two principal components explain 45% of the variability in the data, compared to 39% for the first two principal components of the landmark data. This may indicate that using the curve data to identify principal components as a means of describing the data with reduced dimensionality may be more effective than using the landmark data. The results are not hugely different, however, which is perhaps not surprising, as we are considering the same faces, and the curves are landmark-based.

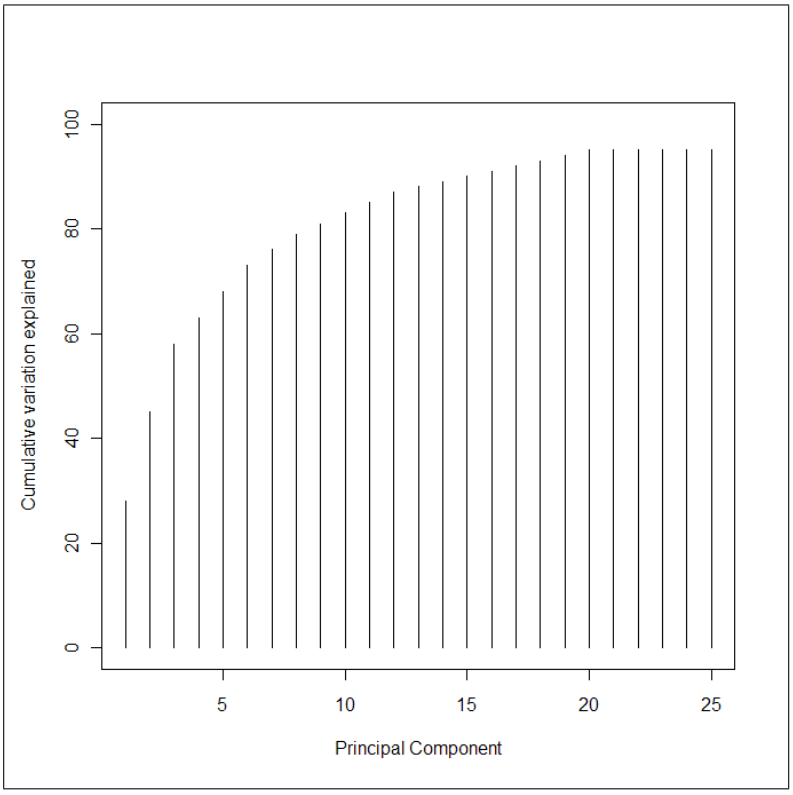

Figure 3.5: Cumulative proportion of the variability in the data explained by each additional principal component.

Figure 3.6 shows boxplots of the first and second principal component scores,

split by sex. The second plot in particular seems to show some differences in the spread of the principal component scores, but a Student's t-test found no significant differences at the 5% level in either the first or second principal component scores between males and females. The majority of subsequent principal component scores also displayed no significant differences between males and females. These results are the same as those that were obtained for the landmark-based principal component analysis.

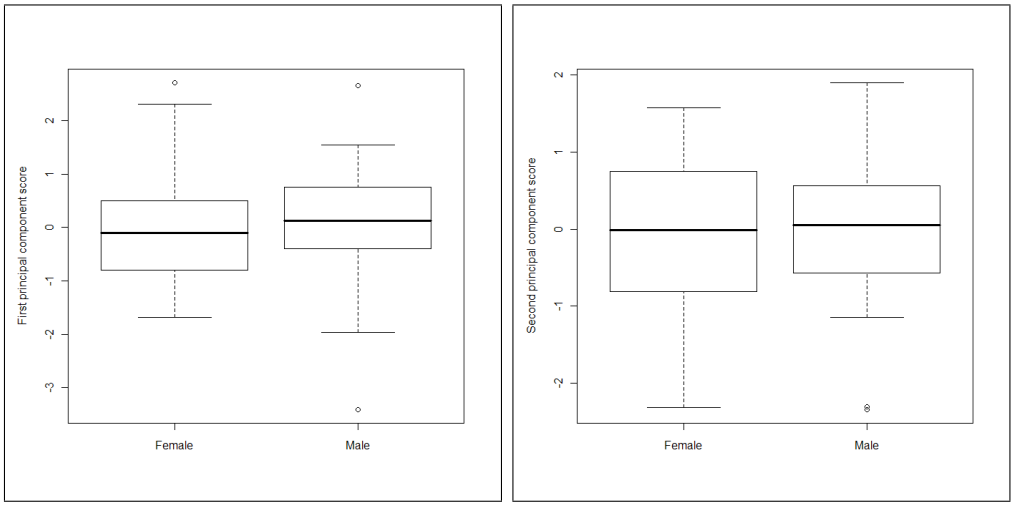

Figure 3.6: Boxplots of the first and second principal component scores, split by sex.

Figure 3.7 plots the first and second principal component scores against one another, split by sex, and, as was the case for the landmark-based analysis, no split is apparent between the sexes.

From comparing the landmark-based and curve-based principal component analyses, we can tentatively conclude that we may see similar results in terms of the analysis of the data, but that the comparative richness of the curve data

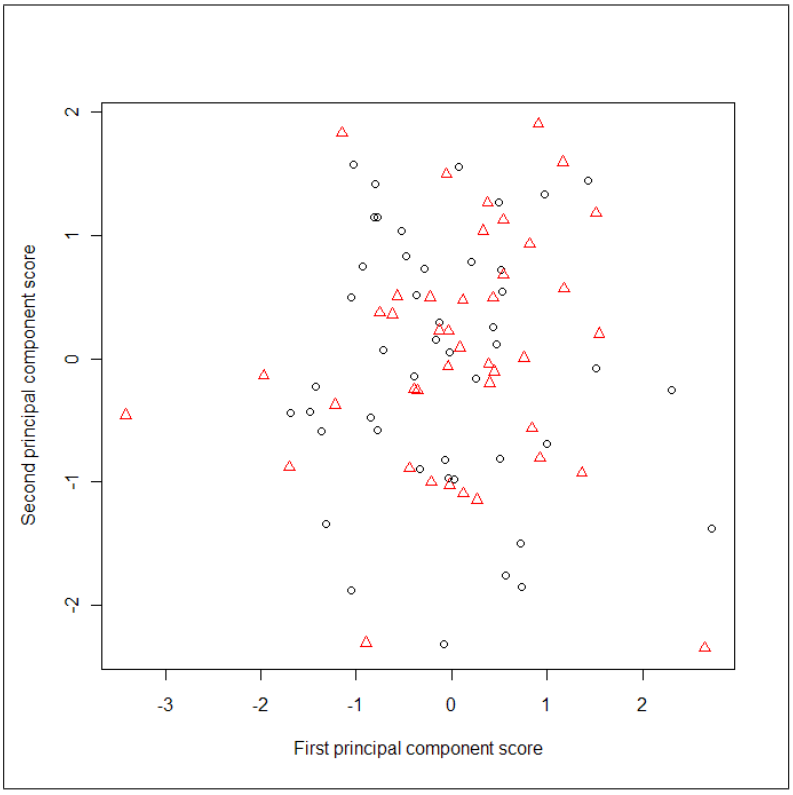

Figure 3.7: Plot of the first against second principal component scores, comparing males (triangles) against females (circles).

may be useful should we wish to employ the technique of dimensionality reduction, which may become progressively more useful as technology allows higher resolution images to be captured.

Chapter 4

A Surface Curvature based Method for Facial Curve Identification

In Chapter 3, a method was discussed and implemented which allowed the identification of a pre-defined set of facial curves. The main characteristics of this method were that each curve had end-points defined by the anatomical facial landmarks, and each curve (or partial curve for later assimilation) was constrained to lie on a plane which was also defined using landmark data. Criticisms of this method were acknowledged at the end of Chapter 3. It is therefore interesting to attempt to use further properties of the face to extract the curves. In this chapter, the degree of curvature at distinct points on the face is used to identify where the facial ridges and valleys occur. Section 4.1 provides a discussion of how the principal curvatures and directions of the surface are calculated and the related definitions of ridges and their corresponding 'red' or 'blue' nature

are then presented in Section 4.2. It has also been necessary to implement the definition of curvature with regards to parameterized curves in two dimensions, so this necessary theory is discussed in Section 4.3. Section 4.4 goes on to outline how the theory behind surface curvature was implemented within R to identify the facial curves of interest.

4.1 Surface Curvature: Principal Curvatures and Principal Directions

The curvature of a surface can be described by the way in which the normal to the surface changes from point to point. Gray *et al.* (2006) provide an excellent discussion of the underlying theory of curvature, and the notation and theory in this section follow from both this text and the summary given by Goldfeather & Interrante (2004). The idea of 'curvature' in three dimensions is best considered through a combination of definitions, each of which provides a different measurement of the way a surface changes around any given point. Curvature functions which are important in describing a surface include the normal curvature k and corresponding principal curvatures k_{max} and k_{min}; the Gaussian curvature K, and mean curvature H, which are all defined at any point $\mathbf{p}$ on a regular surface $\mathcal{M} \in \mathbb{R}^3$.

Given $\mathbf{v_p}$, a tangent vector to $\mathcal{M}$ at $\mathbf{p}$, *normal curvature* is denoted by $k(\mathbf{v_p})$, which is a real number quantifying the extent of bending of $\mathcal{M}$ in the direction $\mathbf{v_p}$. Thus at the point $\mathbf{p}$, there exist many values of normal curvature, each calculated for a different tangent vector to $\mathbf{p}$. Note that all tangent vectors to $\mathbf{p}$ will lie in a single plane, the normal plane to $\mathcal{M}$ at $\mathbf{p}$. Essentially, the values of

normal curvature quantify the degree of bending along the plane curve which is formed by the intersection of the surface $\mathcal{M}$ with the plane passing through the tangent vector $\mathbf{v_p}$, orthogonal to $\mathcal{M}$. The maximum and minimum values of all normal curvatures at $\mathbf{p}$ are known as the *principal curvatures*, denoted k_{max} and k_{min} ($k_{max} \geq k_{min}$), and the tangent vectors at which these extreme values occur are the corresponding *principal directions*.

It can be noted that the Gaussian curvature and mean curvature are related to the principal curvatures by

$$K = k_{max}k_{min} \qquad \text{and} \qquad H = \frac{1}{2}(k_{max} + k_{min}).$$

The Gaussian and principal curvatures allow a distinction to be made between different types of points on a surface, as shown in Table 4.1, for any point $\mathbf{p}$. Note that, with the exception of the umbilic, properly umbilic and planar points, which can be considered as special cases, all point on the surface can be classified as either elliptic convex, hyperbolic (saddle-shaped), or elliptic concave , with the curves separating these areas known as parabolic curves. (Additional information is taken from Kühnel & Hunt (2006) and Hallinan *et al.* (1999).)

The calculation of the principal curvatures and directions is made straightforward through the implementation of some important matrix theory. At any point, $\mathbf{p}$, a linear operator called the *shape operator*, can be considered through its corresponding *Weingarten matrix*, which quantifies the bending of $\mathcal{M}$ at $\mathbf{p}$. Let N_p denote the unit normal to $\mathcal{M}$ at $\mathbf{p}$, and define $X(u, v)$ as a local parametrization of $\mathcal{M}$ in a neighbourhood of $\mathbf{p}$, using $X_u(p)$, $X_v(p)$, N_p as a local coordinate system. Then the Weingarten matrix is defined by:

Type of Points	Gaussian Curvature	Principal Curvatures
Elliptic convex	$K(\mathbf{p}) > 0$	$k_{min}, k_{max} > 0$
Elliptic concave	$K(\mathbf{p}) > 0$	$k_{min}, k_{max} < 0$
Hyperbolic	$K(\mathbf{p}) < 0$	k_{min}, k_{max} have opposite sign
Parabolic	$K(\mathbf{p}) = 0$	Exactly one of k_{min}, k_{max} is zero
Umbilic	$K(\mathbf{p}) \geq 0$	$k_{min} = k_{max}$
Properly Umbilic	$K(\mathbf{p}) > 0$	$k_{min} = k_{max} \neq 0$
Planar	$K(\mathbf{p}) = 0$	$k_{min} = k_{max} = 0$

Table 4.1: Definitions of different types of surface point, based on Gaussian and Principal curvatures

$$W(\mathbf{p}) = \frac{-1}{EG - F^2} \begin{pmatrix} Ge - Ff & Gf - Fg \\ -Fe + Ef & -Ff + Eg \end{pmatrix}$$

where

$$\begin{aligned} e &= N_p \cdot X_{uu}(p) & E &= X_u(p) \cdot X_u(p) \\ f &= N_p \cdot X_{uv}(p) & F &= X_u(p) \cdot X_v(p) \\ g &= N_p \cdot X_{vv}(p) & G &= X_v(p) \cdot X_v(p). \end{aligned}$$

The eigenvalues of the Weingarten matrix $W(\mathbf{p})$ of the shape operator S at $\mathbf{p} \in \mathcal{M}$ are then the principal curvatures k_{max} and k_{min} of $\mathcal{M}$ at $\mathbf{p}$. The corresponding unit eigenvectors of W are the unit principal vectors $\mathbf{e_1}$ and $\mathbf{e_2}$ and the principal directions $\mathbf{t}_{max}$ and $\mathbf{t}_{min}$ arise from these vectors. Equally, the Gaussian curvature can be defined as the determinant of the shape operator, and

the Mean curvature as half the trace of the shape operator, i.e.

$$K(\mathbf{p}) = \det(W(\mathbf{p})) \qquad \text{and} \qquad H(\mathbf{p}) = \frac{1}{2}\text{tr}(W(\mathbf{p})).$$

Another way to interpret the principal curvatures of a point on a surface is to consider the *Monge form* of the surface in a local neighbourhood of the point. The Monge form of the local surface, or *Monge patch* is of the form:

$$\begin{aligned} z = f(x,y) &= \frac{1}{2}(k_{max}x^2 + k_{min}y^2) + \frac{1}{6}(b_0x^3 + 3b_1x^2y + 3b_2xy^2 + b_3y^3) + \\ &\quad \frac{1}{24}(c_0x^4 + 4c_1x^3y + 6c_2x^2y^2 + 4c_3xy^3 + c_4y^4) + \text{higher order terms} \end{aligned}$$

where the coordinate system is defined so that the tangent plane to the point is $z = 0$, and the surface is then rotated in the x-y plane so that the x- and y- axes are the principal directions. Then k_{max} and k_{min} are the principal curvatures.

4.2 Ridges

Consider a polyhedral surface, with 'edges' where two faces meet. It is relatively straightforward to both visually recognise where these edges occur and also to define them geometrically. As an example, consider the unit cube in Figure 4.1. It is obvious that the edges of this cube (the blue and red lines) can be identified easily by a visual examination, and it is also a simple matter to mathematically define these edges, for example the red edge shown is defined by

$$x = 0 \ , \ y = 0 \ , \ 0 \leq z \leq 1.$$

It is desirable to have a corresponding definition for non-polyhedral surfaces, in particular smooth surfaces in $\mathbb{R}^3$, such as a face, so that curves and features

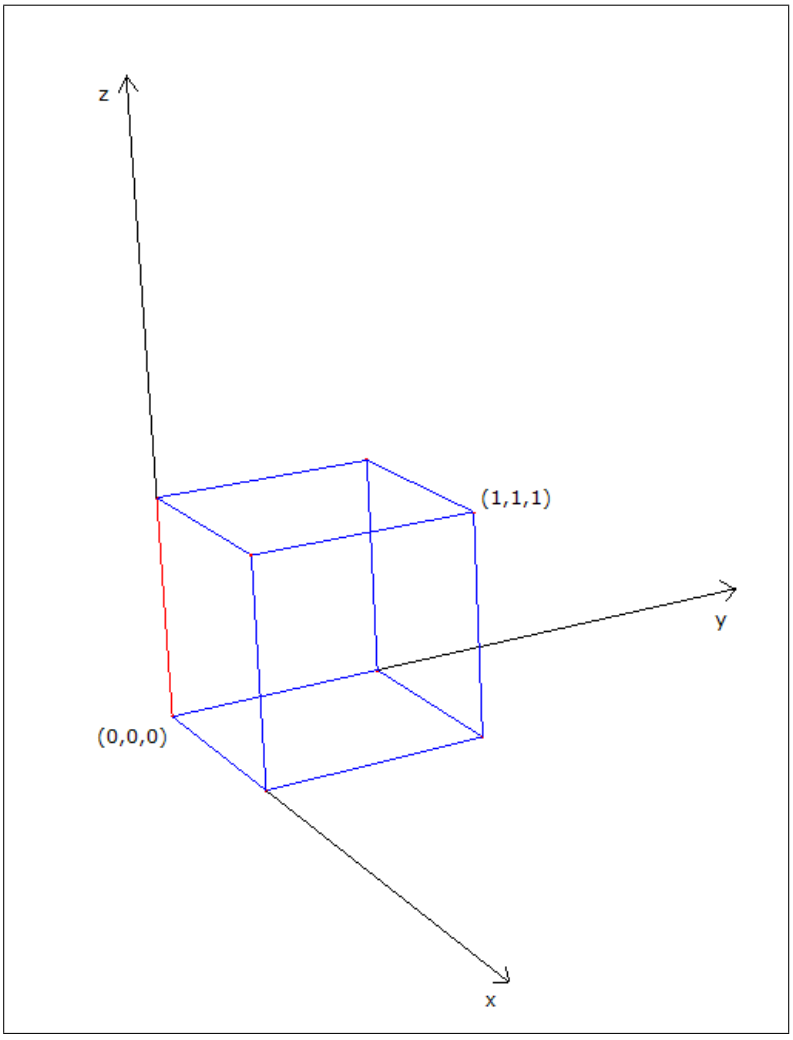

Figure 4.1: Unit cube

inherent to the face can be defined, detected and studied. The features which describe features of the surface in the same way that edges do for a polyhedral shape are known as *ridges* and can be subdivided into either 'blue' or 'red' ridges. Blue ridges are those which would be the extension of convex polyhedral edges, and red ridges the extension of concave polyhedral edges. Together these provide an important picture of a surface in their own right.

Since ridge lines mark the main geometrical boundaries on a surface, it is perhaps no surprise that they are defined using principal curvature and principal direction information, and the outline of the theory to follow is adapted from Hallinan *et al.* (1999), Ohtake *et al.* (2004) and Page *et al.* (2006).

Ridge lines are essentially curves on a surface along which the surface bends sharply, composed of a series of *ridge points*. Blue ridge points can be defined by the local maximum value of k_{max} along the maximum principal direction curvature line, while red ridge points are defined by the local minimum value of k_{min} along the minimum principal direction curvature line.

Let $\mathcal{M}$ represent the smooth surface in which we are interested, with principal curvatures k_{max} and k_{min} at a point $\mathbf{p} \in \mathcal{M}$, and corresponding principal directions $\mathbf{t}_{max}$ and $\mathbf{t}_{min}$. Now define the derivatives of the principal curvatures along their associated directions as:

$$c_{max} = \frac{\partial k_{max}}{\partial \mathbf{t}_{max}} \qquad \text{and} \qquad c_{min} = \frac{\partial k_{min}}{\partial \mathbf{t}_{min}}.$$

These derivatives allow the identification of principal curvature maxima and minima via the zero-crossings of c_{max} and c_{min}. So geometrically, blue ridges are characterized by the following conditions:

$$c_{max} = 0, \qquad \frac{\partial^2 k_{max}}{\partial \mathbf{t}_{max}{}^2} = \frac{\partial c_{max}}{\partial \mathbf{t}_{max}} < 0, \qquad k_{max} > |k_{min}|.$$

That is, the point of interest is a local maximum turning point in the maximum principal curvature direction (from the first two conditions), where the magnitude of the degree of curvature is greatest in the maximum principal curvature direction (from the third condition). Red ridges are similarly characterized by:

$$c_{min} = 0, \qquad \frac{\partial^2 k_{min}}{\partial \mathbf{t}_{min}{}^2} = \frac{\partial c_{min}}{\partial \mathbf{t}_{min}} > 0, \qquad k_{min} < -|k_{max}|.$$

Here, the point of interest is a local minimum turning point in the minimum principal curvature direction, where the magnitude of the degree of curvature is greatest in the minimum principal curvature direction. In terms of the surface in Monge form (as discussed in Section 4.1), blue ridge points are characterized by the disappearance of the x^3 term, that is, $b_0 = 0$, while red ridge points are characterized by $b_3 = 0$, i.e. the disappearance of the y^3 term. Ridges can also be further subdivided into 'elliptic' or 'hyperbolic', however this shall not be discussed here. A detailed discussion of ridges, and related surface features is presented in Hallinan *et al.* (1999).

4.3 Curvature of Parametrized Plane Curves

The algorithm used in R to identify the ridge lines discussed in the previous section is outlined in 4.4, and involves some analysis of plane curves, which are curves which lie in a single two-dimensional plane. It is therefore necessary to introduce the theory behind the analysis of curvature of plane curves. The relevant theory, taken from Gray *et al.* (2006) is therefore discussed in this section.

At any point, the curvature of a curve essentially quantifies the extent to which the curve fails to be a straight line. Consider a parameterized curve $\boldsymbol{\alpha}(t) =$

$(x(t), y(t))$. The curvature $\boldsymbol{\kappa 2}[\boldsymbol{\alpha}](t)$ of $\boldsymbol{\alpha}(t)$, where the "2" represents the fact that the curve is in 2 dimensions, is given by

$$\boldsymbol{\kappa 2}[\boldsymbol{\alpha}](t) = \frac{\dot{x}\ddot{y} - \ddot{x}\dot{y}}{(\dot{x}^2 + \dot{y}^2)^{3/2}}$$

where the derivatives are with respect to t. The value of the curvature is independent of the parametrization up to sign. Thus to evaluate the degree of curvature of a plane curve at any given point, it is necessary to calculate both first and second derivatives of the x and y positions of the curve independently, with respect to the position of the point along the length of the curve.

4.4 Surface Curvature estimation for Facial Mesh Data

At any point $\mathbf{p}$ on the facial mesh surface, it is firstly necessary to estimate the Weingarten matrix, W, as defined in section 4.1. Using the local coordinate system which was introduced in section 4.1, consider points in a local neighbourhood of $\mathbf{p}$, and fit the surface

$$f(x, y) = \frac{A}{2}x^2 + Bxy + \frac{C}{2}y^2 + Dx^3 + Ex^2y + Fxy^2 + Gy^3 \qquad (4.1)$$

to the points using least-squares, as discussed in Goldfeather & Interrante (2004). By fitting this cubic surface, which is more complex than, for example, a purely quadratic surface, Goldfeather & Interrante (2004) found that errors in principal direction estimation were reduced. The estimated Weingarten matrix for this surface is

$$W' = \begin{pmatrix} A & B \\ B & C \end{pmatrix}$$

and this can be used as an approximation for the true Weingarten matrix of the surface, with principal curvatures and directions estimated through eigenanalysis of W'.

Ohtake *et al.* (2004) implemented an effective method for identifying ridge points using the derivative-based criteria, which determines whether each point on the mesh is, or is not, a ridge point. Since fitting a surface to estimate the principal curvature and directions for every mesh point on every face would be very computationally intensive, an alternative method is used here. The aims of the method are twofold:

1. To identify facial curves corresponding to those in the previous chapter (i.e. with endpoints defined by landmarks), but with the flexibility to deviate from a planar surface.

2. To establish whether any interesting ridge curves appear on the face, and if so, to consider whether these ridge curves provide any information about the surface unquantified by the previously extracted curves.

It seems reasonable to assume that the desired facial curves (illustrated in Figure 3.4) will not deviate hugely from the planar curves, so to reduce computational time, it was decided to use the extracted planar points, as shown in Figure 4.2 for one subject, as starting points for identifying local ridge points.

For each planar point, $\mathbf{p}$, a set of local axes was first defined so that the local z-axis was an estimate of the normal to the surface at $\mathbf{p}$, which was estimated as the normalized average of the unit normal vectors arising from all adjacent triangular mesh segments. Methods to estimate a more accurate surface normal include various weighted averages, however Goldfeather & Interrante (2004) conclude that

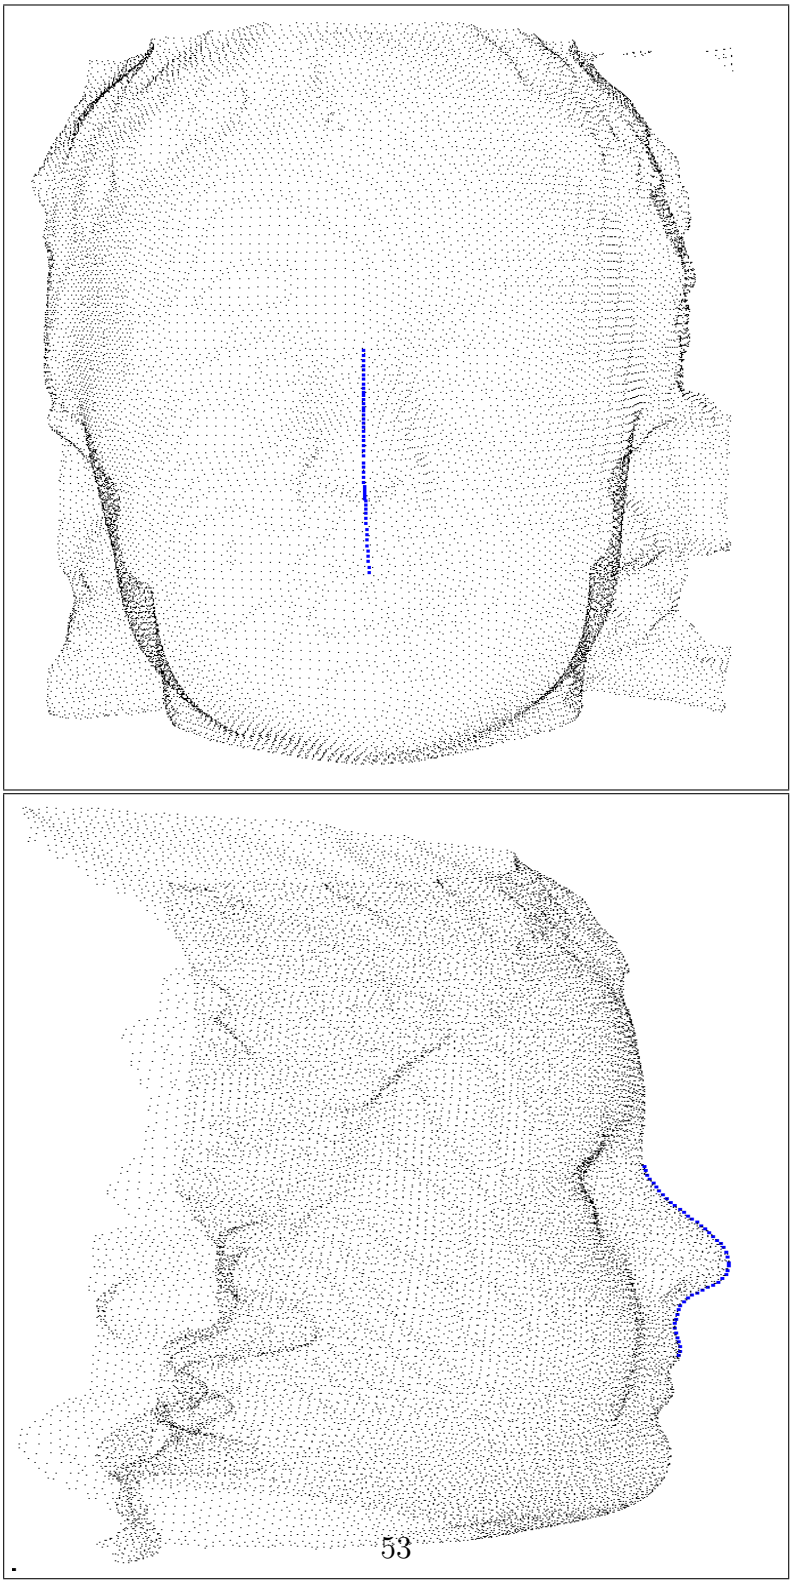

Figure 4.2: Midline Curve for one subject.

the differences in surface normal which are estimated using the different methods are not significant in practice, therefore in this case, the unweighted method was used. The x- and y- axes were chosen as arbitrary vectors, orthogonal to both the z-axis and each other. As an example, consider the set of local axes for a single point on the midline plane shown in Figure 4.3. This point will be used to illustrate concepts for the remainder of this chapter.

The surface in Equation (4.4) was then to be fitted to the local area. It was important to choose an area of a size which captured the behaviour of the surface around $\mathbf{p}$ adequately, without selecting an area so large that the closest fitted surface would not capture the way the surface changed close to $\mathbf{p}$. Various sizes were considered for local area, ranging from 3mm to 3cm radii from $\mathbf{p}$. Both the smaller and larger values led to surfaces which, when viewed alongside the surface mesh, did not capture the characteristics of the local area sufficiently. The decision was therefore made to consider, as a local area, those points which lay within 1cm of the point $\mathbf{p}$. This is by no means the only acceptable value, and it should also be observed that while this value of 1cm seems to perform well for the meshes representing 5-year-old children, and for the area around the midline curve in particular, it is anticipated that this would need to be reconsidered for meshes with different characteristics. For example, we would anticipate that the facial characteristics of younger children would occur closer together, and similarly we would expect that the corresponding characteristics of older children and adults would occur further apart. An illustration of an area to be used for surface fitting is shown in Figure 4.4.

To fit the surface, it was first necessary to transform the selected points into local coordinates, so for the local coordinate system ($x = (x_1, x_2, x_3), y =$

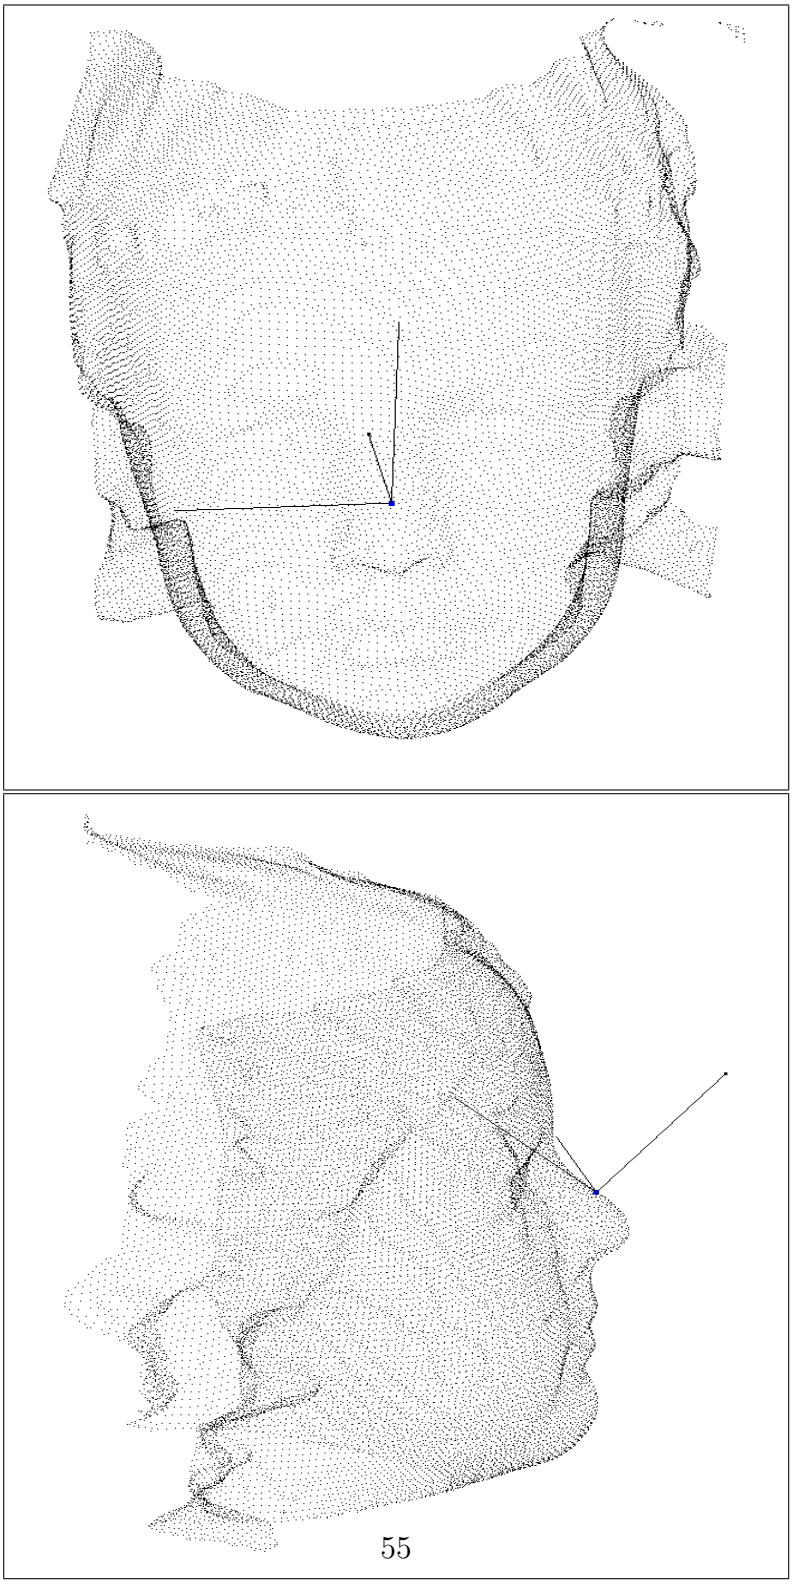

Figure 4.3: Local axes for one point on the planar midline curve, with the z-axis depicted by the axis line with a point on its end.

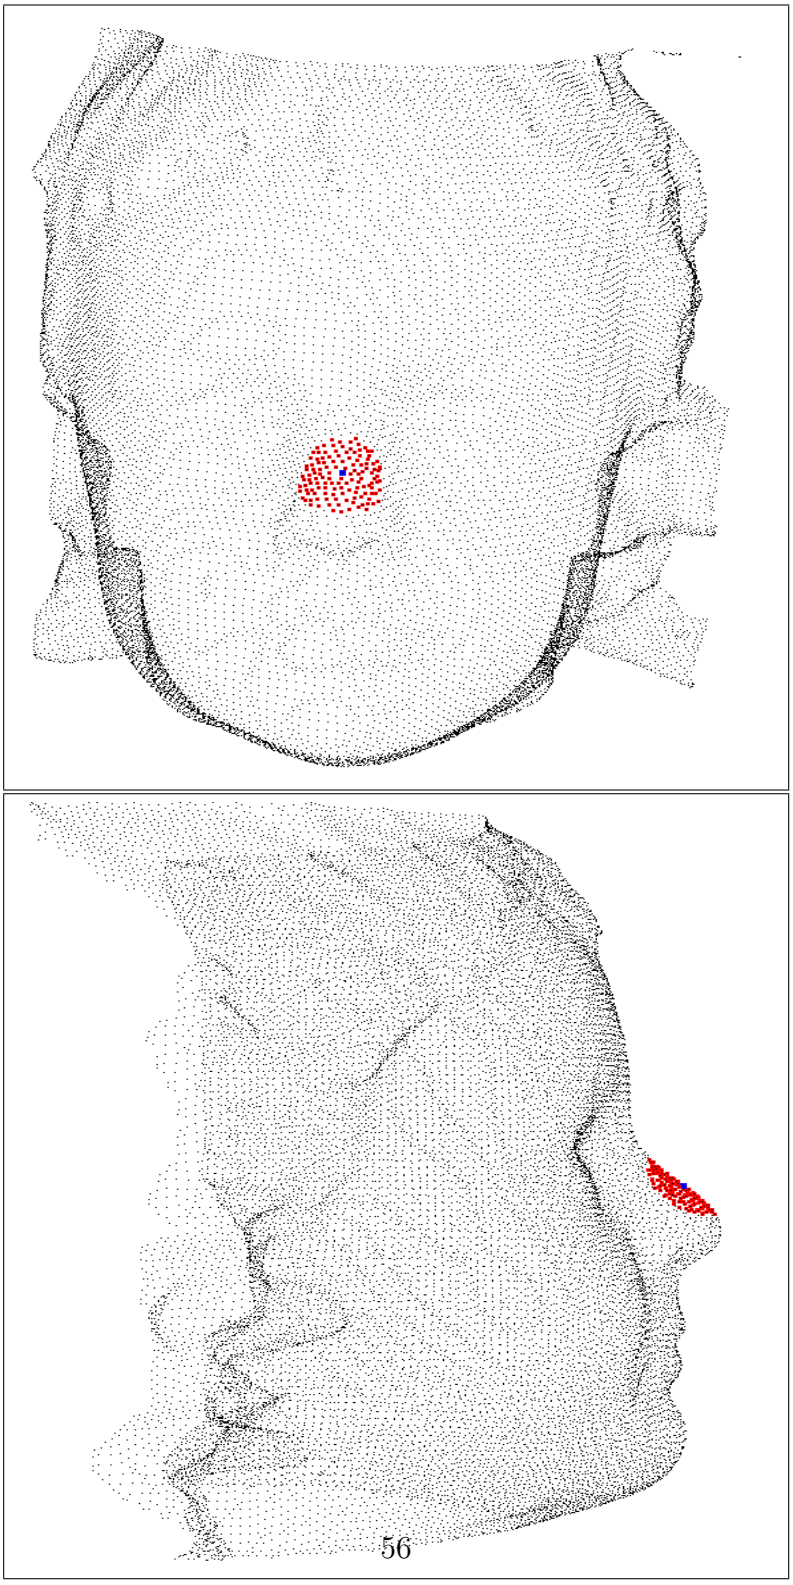

Figure 4.4: Local area for surface fitting (shown in red) around one point (shown in blue).

$(y_1, y_2, y_3), z = (z_1, z_2, z_3))$, with point of origin $\mathbf{p}$ and each of x, y and z a unit vector, the coordinate transformation will represent $\mathbf{p}$ as $(0, 0, 0)$, and each local point $\mathbf{a} = (a_1, a_2, a_3)$ will be represented as:

$$((a_1, a_2, a_3) - (p_1, p_2, p_3)) \begin{pmatrix} x_1 & y_1 & z_1 \\ x_2 & y_2 & z_2 \\ x_3 & y_3 & z_3 \end{pmatrix}$$

The surface in Equation 4.4 was then fitted to the transformed chosen points using least squares, and the Weingarten matrix

$$W' = \begin{pmatrix} A & B \\ B & C \end{pmatrix}$$

for this surface was then used as an estimate for the true Weingarten matrix of the mesh surface. The principal curvatures were then estimated by the eigenvalues of W', with the corresponding eigenvectors estimating the principal directions. The maximum principal curvatures for two faces are shown in Figure 4.5, while the minimum principal curvatures for the same two faces are shown in Figure 4.6.

It is clear that certain areas of the face, such as the lips and nose are distinct in their principal curvature characteristics, and also that there is a small degree of overlap in the values for maximum and minimum principal curvatures. It is also interesting to consider a visual representation of the Gaussian curvature $(= k_{max}k_{min})$ which illustrates which areas of the face are elliptic convex ($k_{max} > k_{min} > 0$, areas shown in green), elliptic concave ($k_{min} < k_{max} < 0$, shown in red), or hyperbolic ($k_{min} < 0 < k_{max}$, shown in blue). This can be seen in Figure 4.7. Note that the characteristics of these areas may at first seem to be in reverse, for example, one would have expected the sunken eye areas to have

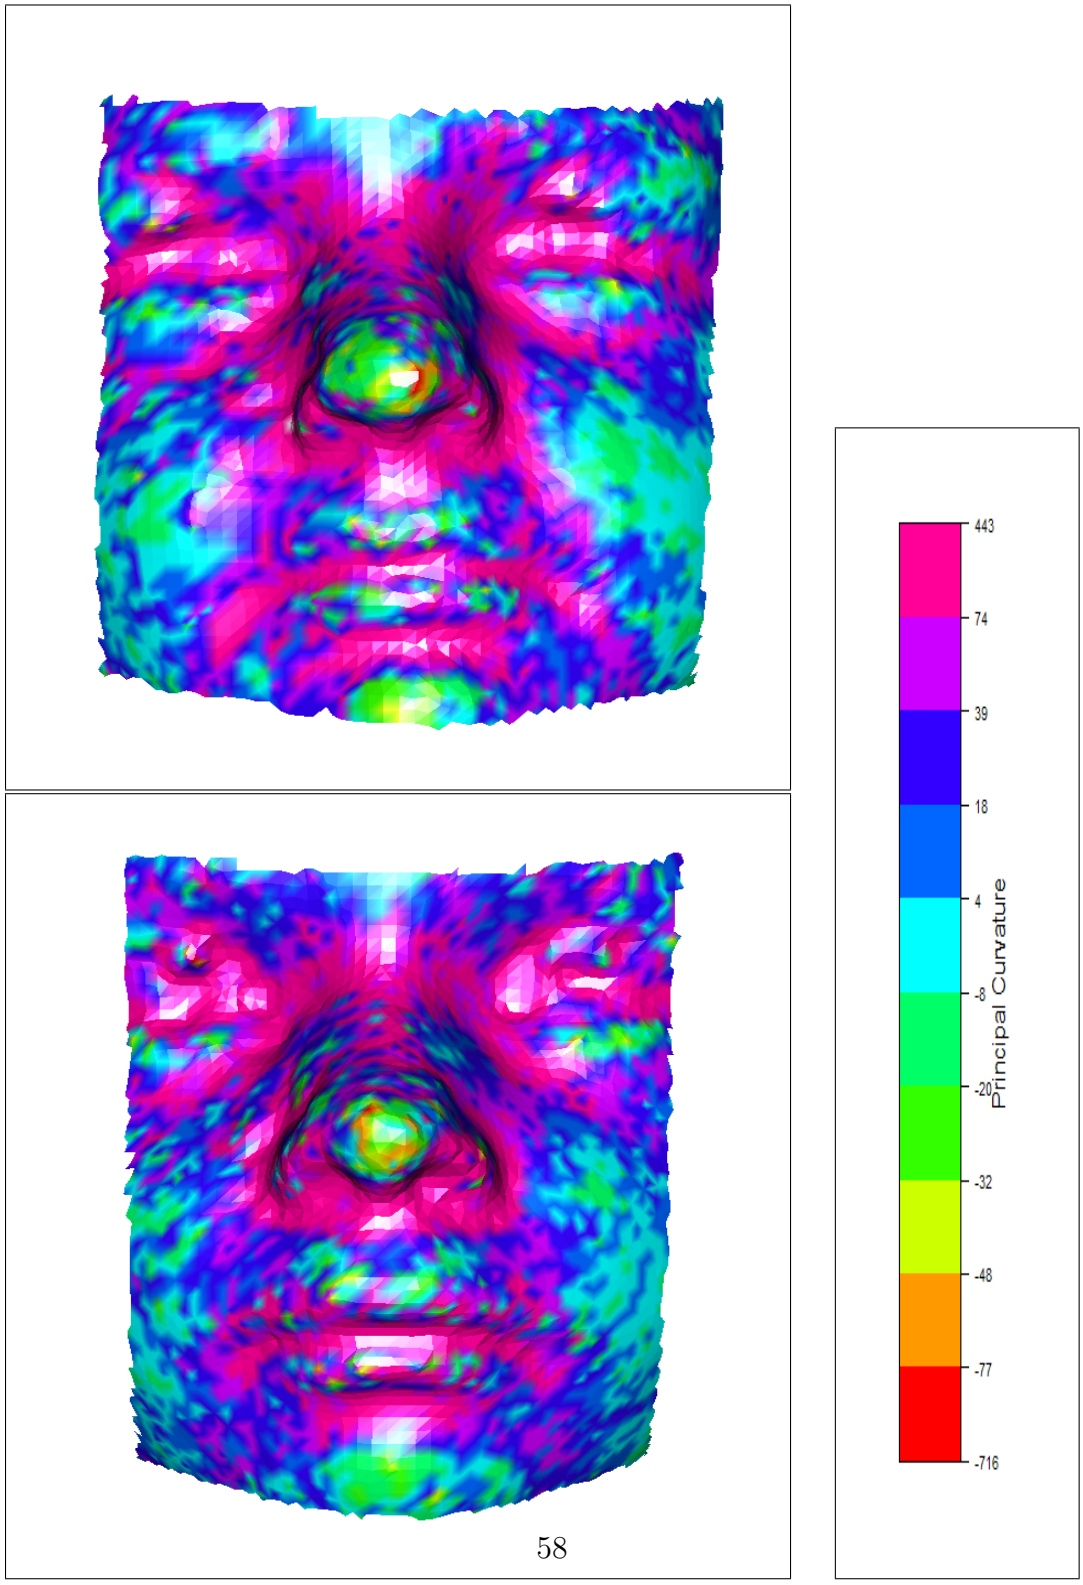

Figure 4.5: Values of Maximum Principal Curvature for two faces.

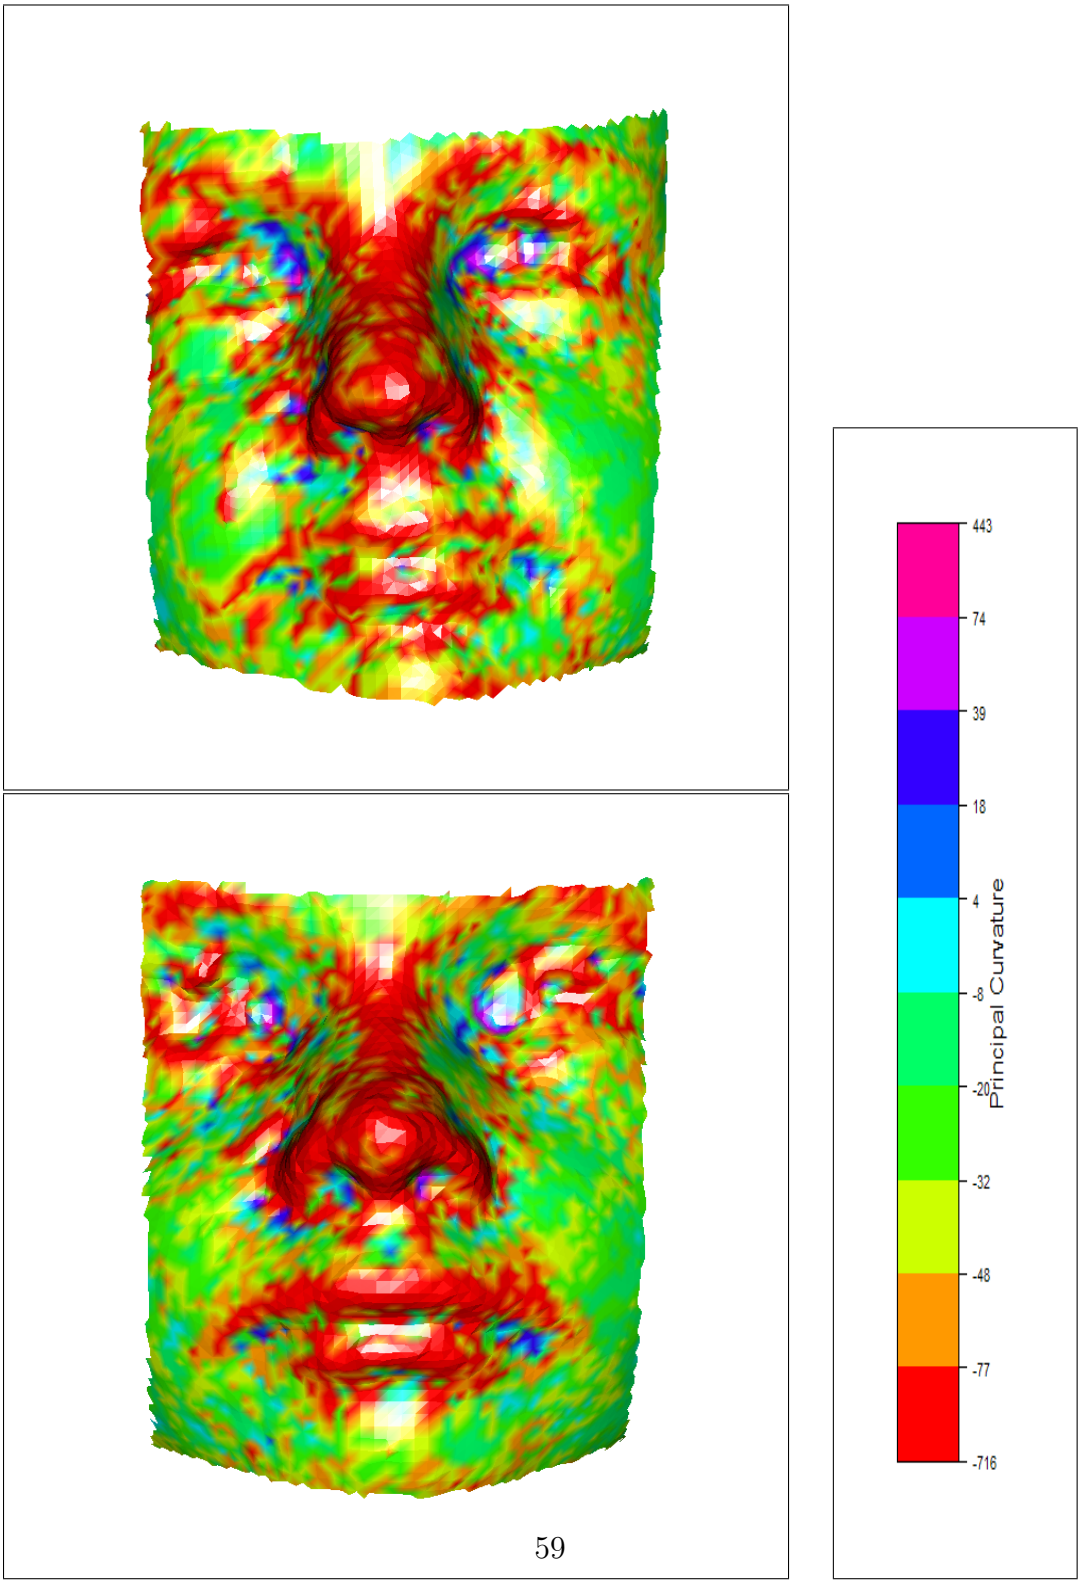

Figure 4.6: Values of Minimum Principal Curvature for two faces.

concave characteristics, and the protruding nasal area to be convex, but it should be remembered that the gaussian curvature is defined in relation to the normal to the face at a given point. Thus although a point at the tip of the nose may seem intuitively convex, if one recalls that the given point is the origin, with z-axis defined by the normal which points from the face towards the viewer, then it is more obvious that the surface is, in this context, concave.

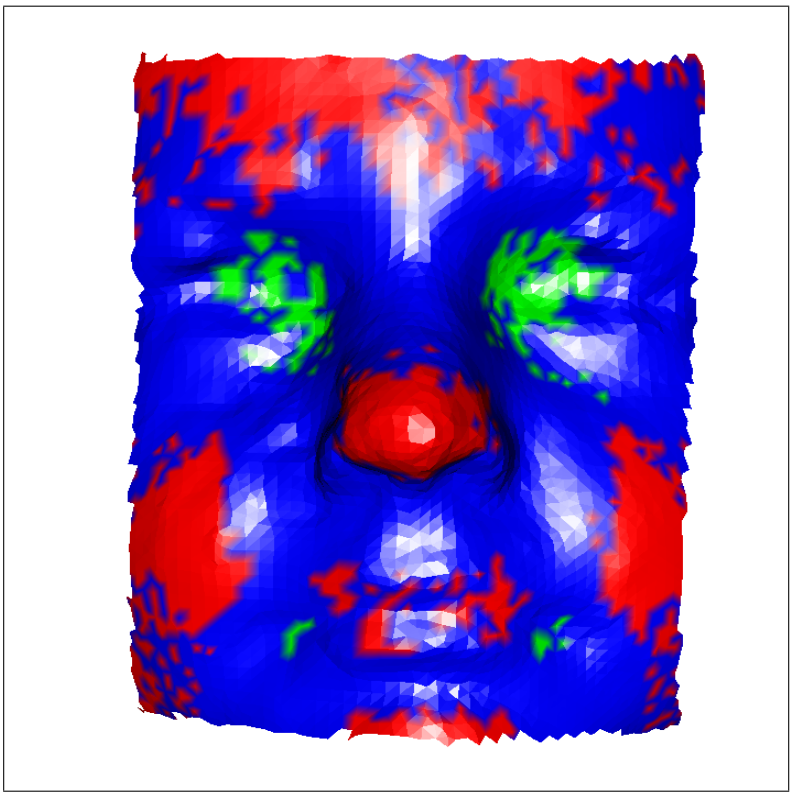

Figure 4.7: Gaussian curvature for each mesh point on a selected area of one face, with elliptic convex areas shown in green, elliptic concave areas shown in red and hyperbolic areas shown in blue.

The estimated principal directions at each mesh point also provide useful information about the surface behaviour. A selected area of the face is shown in Figure 4.8, depicting maximum estimated principal curvature directions and minimum estimated principal curvature directions. Both are shown on the same image in Figure 4.9, and the necessary orthogonal nature of the relationship between maximum and minimum principal directions is obvious.

The way in which the surface principal directions change is apparent, although there does appear to be a certain degree of variability in the directions, for example in the area around the tip of the nose. This may be due to the fact that the tip of the nose is an area where the normal curvatures are very close to each other, i.e. the surface is close to being an umbilic. At an umbilic point, the normal curvatures are equal for every direction, and so there are no maximum or minimum principal curvatures or related directions. The result of this is that even at a point which is close to being an umbilic, the errors in estimating the direction will be noticeable. It should therefore be kept in mind that the potential errors in principal direction estimation may affect the further techniques to be implemented.

4.5 Identifying Ridge Points on the Facial Mesh

The previous section described the mechanism by which, at each required point, the principal curvatures and directions may be calculated. This section implements a method which uses this information at each previously determined planar point to identify a local ridge point. It is of interest to investigate whether any ridge curves correspond to the facial curves of interest, in particular the

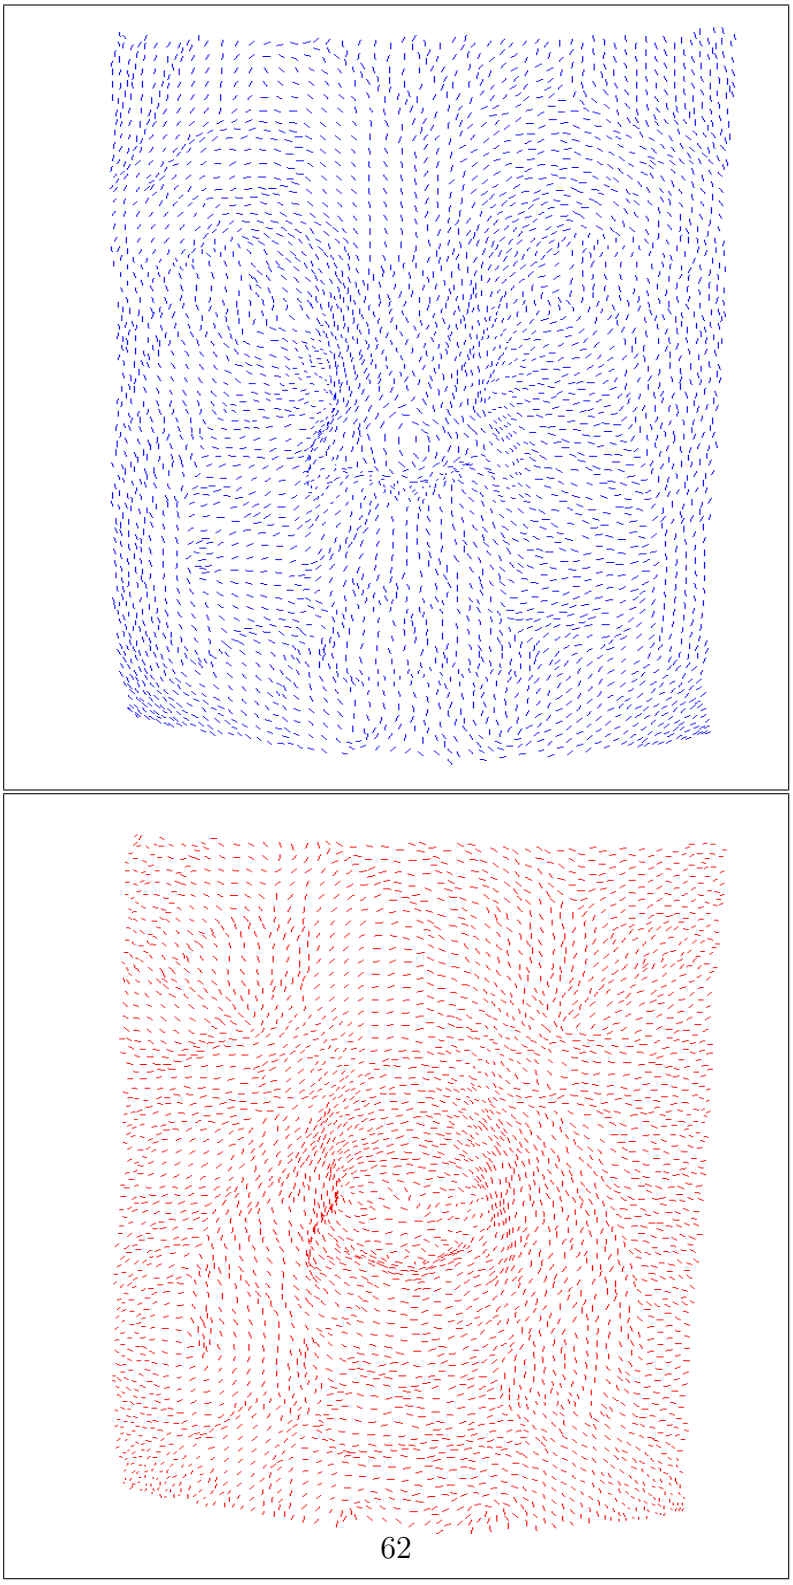

Figure 4.8: Maximum (Blue) and Minimum (Red) estimated Principal Directions shown for each mesh point on a selected area of one face.

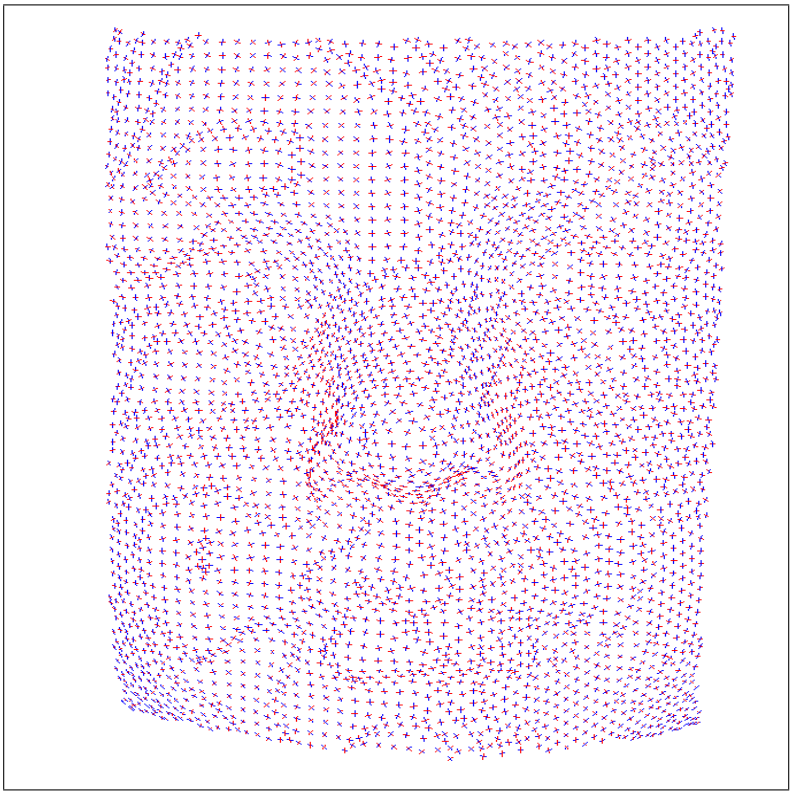

Figure 4.9: Maximum (Blue) and Minimum (Red) estimated Principal Directions shown together for each mesh point on a selected area of one face.

midline curve. Since a number of techniques are combined in this section, all techniques will be illustrated using the midline curve of the face as an example. The main result on which this is based is that a ridge point can be defined by a point of local maximum curvature along the direction of maximum curvature. In this context, because the points are locally defined with the z-axis as the normal to the surface at the point of consideration $\mathbf{p}$, we are in fact interested in those curvature directions which have the greatest degree of bending *away* from the area $z > 0$. Thus we shall be looking for the points of greatest curvature along the *minimum* principal curvature direction.

Consider, as an illustration, the planar curve point which is shown in Figure 4.10, along with the minimum principal curvature direction at this point, extending away from the point on either side (shown in red). The plane which passes through this line of direction, orthogonal to the face, can then be defined and used to identify the points which lie on the mesh along this direction. To consider only a local area, it was decided to allow a ridge point within 2cm of the planar curve to be identified. These local mesh points along the line of minimum curvature are therefore shown in blue.

A new change of coordinate system can then be applied to force this series of points to lie on the plane defined by $z = 0$. Essentially, the three-dimensional problem of identifying the point of maximum curvature along a particular direction has now been reduced to a two-dimensional problem of identifying the point of maximum curvature of a series of points making up a local plane curve. The problem is to identify the point of maximum curvature in the curve which follows the points shown in Figure 4.11. It is at this stage in the method that the plane curve theory discussed in Section 4.3 is implemented.

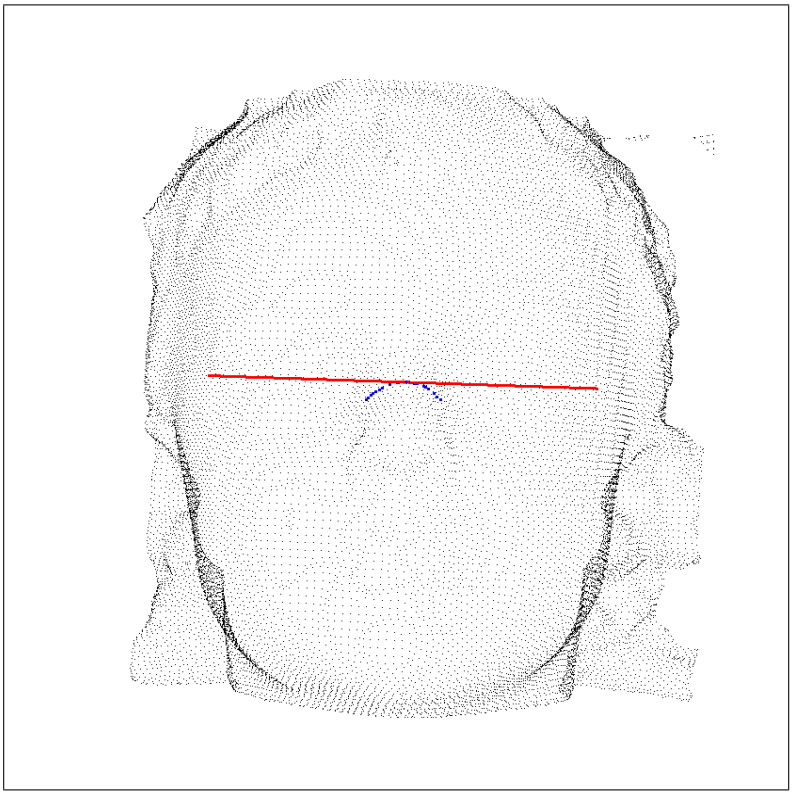

Figure 4.10: Minimum principal curvature direction at a single mesh point, with local mesh points in this direction.

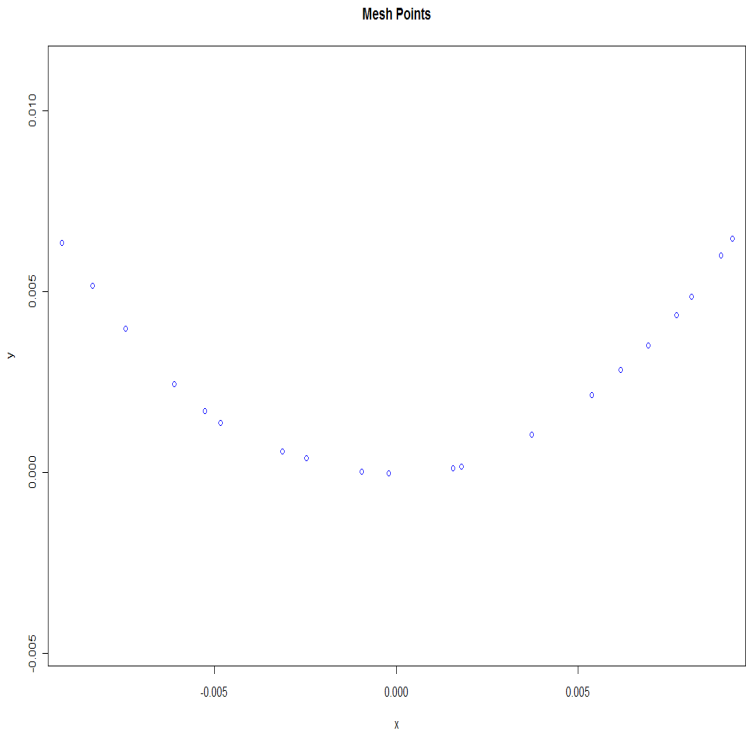

Figure 4.11: Mesh points which lie on the minimum principal direction line, reduced to two dimensions.

It was decided to firstly fit a cubic smoothing spline to the points, as previously described in Chapter 3, and discussed in more detail in Hastie & Tibshirani (1996). We can fit the cubic smoothing spline using the smooth.spline function in R, with the required number of degrees of freedom specified. The function will then efficiently fit a spline to the points making up the facial curve and return the required number of (evenly spaced out on t) predicted points along the length of the spline. This has the effect of representing the planar curve by two separate curves, $x(t)$ and $y(t)$, with t indexing the length of the curve. The first and second derivatives of the spline at each point (and hence the estimated first and second derivatives of the facial curve) are also returned by the function. The first derivatives are denoted by $\dot{x}$ and $\dot{y}$ respectively, where $\dot{x} = \frac{dx}{dt}$ and $\dot{y} = \frac{dy}{dt}$. The second derivatives are correspondingly defined. It should be noted that the valued of the derivatives are not affected by the orientation of the curve. In order for the spline to fit the facial curve fairly well, a relatively high number of degrees of freedom, of 8, is required. 8 is sufficient to describe the curve well because the surface is localised and already fairly smooth. The curvature at each predicted point was then estimated by:

$$\boldsymbol{\kappa 2}[\boldsymbol{\alpha}](t) = \frac{\dot{x}\ddot{y} - \ddot{x}\dot{y}}{(\dot{x}^2 + \dot{y}^2)^{3/2}}$$

where t represents the distance along the curve, as discussed in Section 4.3. It can be seen that the orientation of the curve does not affect the estimated curvature at each point. Figure 4.12 shows an example of a series of predicted points, with the estimated curvature at each point.

Because the orientation of the curve in two-dimensions is not fixed, and may vary depending on the coordinated transformation necessary to force the points onto the $z = 0$ plane, the series of points may take one of many orientations.

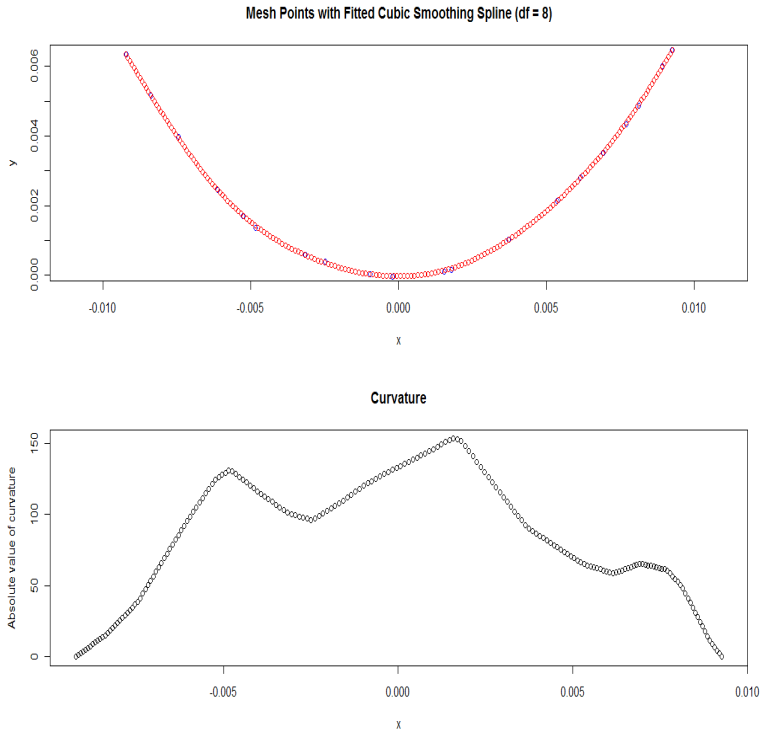

Figure 4.12: Predicted points for a cubic smoothing spline with 8 degrees of freedom, with corresponding absolute value of curvature estimates.

It is therefore not predictable whether the point we seek of maximum curvature will have a high positive, or high negative value. Therefore we consider the point along the predicted curve which has a maximum absolute value of curvature. This point is identified for the illustrated predicted points in Figure 4.13.

However, it is immediately apparent that this is not the expected point of maximum curvature. This problem is a result of the high degrees of freedom allowed to the spline, which seems to lead to first and second derivative estimates which are highly variable. Some experimentation with various degrees of freedom

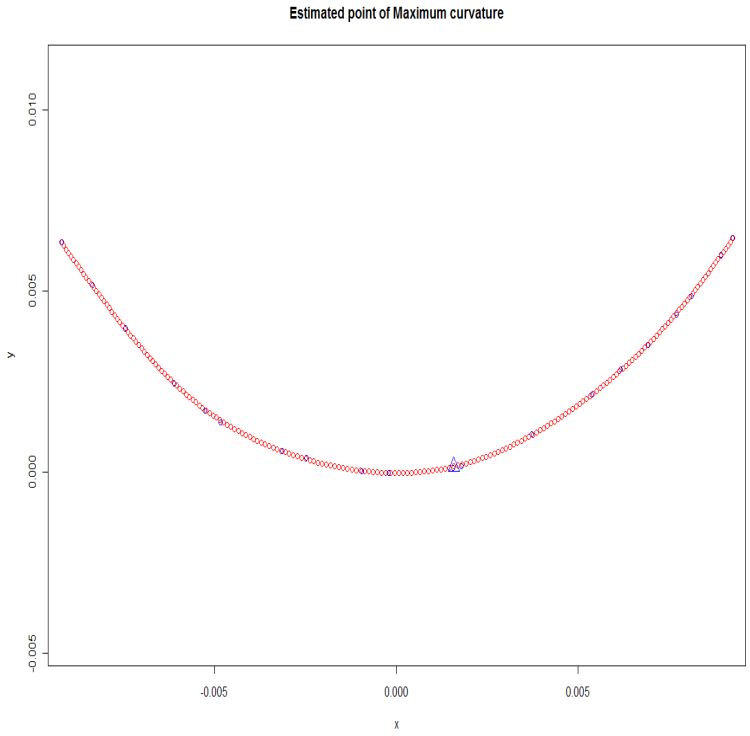

Figure 4.13: Predicted points for a cubic smoothing spline with 8 degrees of freedom, with point of maximum curvature estimated.

leads to the conclusion that using a smaller value gives a much less well-fitting spline, but on this curve the point of maximum curvature is much more reliably identified, as in Figure 4.14.

Therefore it was decided to combine the two prediction methods as follows:

1. Fit a cubic smoothing spline with 4 degrees of freedom to the mesh points. This curve will be a reasonably poor fit.

2. Identify the point of maximum absolute value of curvature for this poorly fitting spline. This will be well estimated

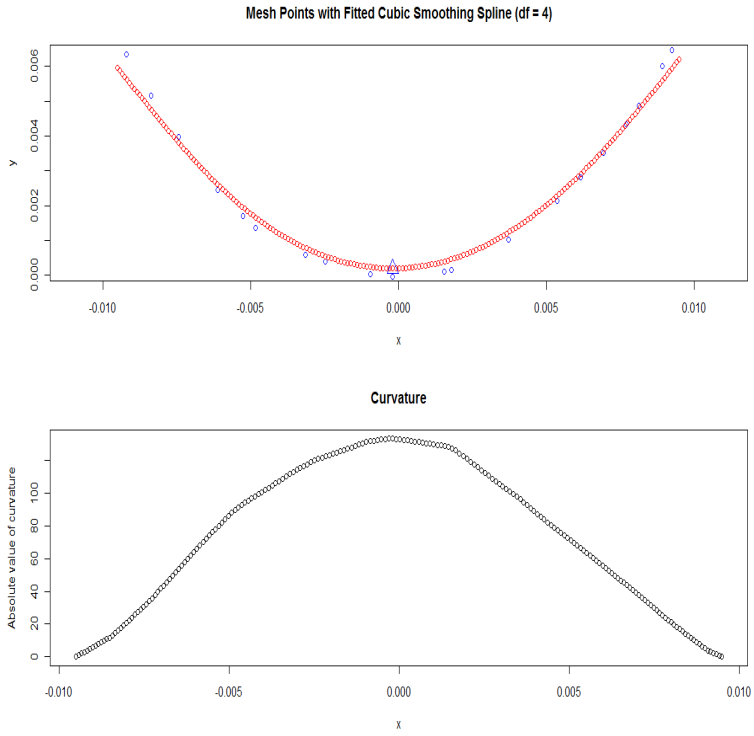

Figure 4.14: Predicted points for a cubic smoothing spline with 4 degrees of freedom, with point of maximum curvature estimated and corresponding curvature estimates.

3. Fit a cubic smoothing spline with 8 degrees of freedom to the points. This should be a good fit.

4. Identify the point on this series of predicted points which corresponds (in terms of distance along the curve) to the point of maximum curvature on the smoothing spline with lower degrees of freedom.

The effect of these steps is shown in Figure 4.15, so the result is that the point of maximum curvature on the $z = 0$ local coordinates plane is identified.

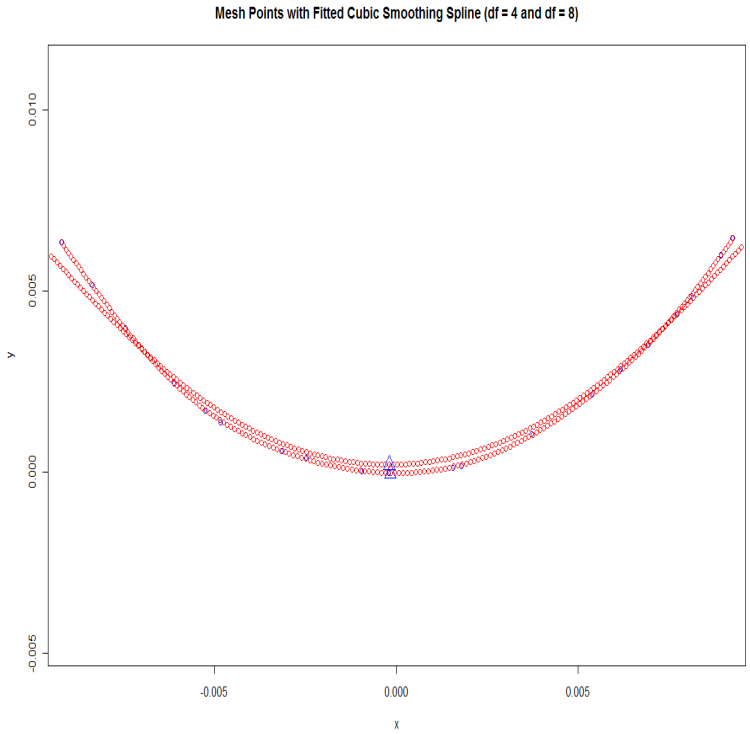

Figure 4.15: Using a combination of two splines to identify the point of maximum curvature.

Transforming back to the original coordinate system then gives the ridge point closest to the planar starting point. Figure 4.16 shows all of the ridge points (blue) which are identified using the planar curve points (red) as starting points.

It is clear that these do not, even when smoothed sufficiently, produce a midline curve which corresponds to the planar curve. At this point, the intention is to use the theory of ridge points to provide an alternative method to planar curve extraction for identifying the desired facial curves. It seems that the midline curve will not be well defined by a series of ridge points with no additional

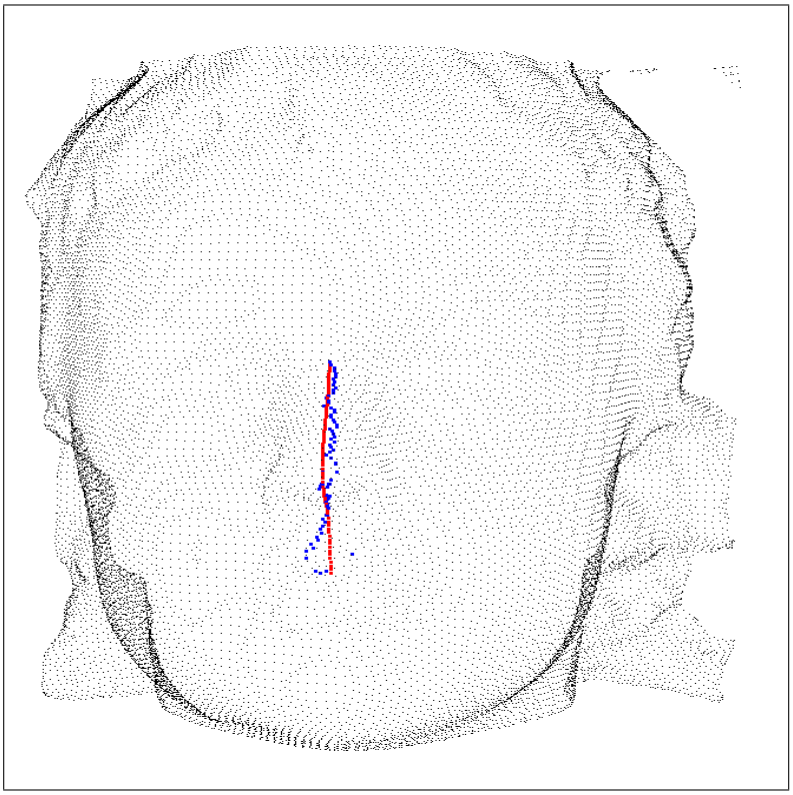

Figure 4.16: Ridge points identified from starting points of the planar midline curve.

constraints. By limiting the points to be used in defining the curve to only ridge points, which may well be good estimates of true ridge points, but which potentially deviate hugely from the original planar curves, we are imposing an unnecessary constraint. As an illustration, it can be seen that there are noticeable deviations from the midline plane around the philtral area in particular.

A proposed refinement of this method involves, rather than the minimum principal curvature direction, the use of, for every point on the planar curve, the direction defined by the vector which runs 'horizontally', from one inner eye

landmark to the other ("enR" to "enL"), as shown in Figure 4.17. This refinement makes use of the knowledge we have of the face, and helps avoid the problems that occur due to variations in the facial surface, which can produce misleading principal curvature directions.

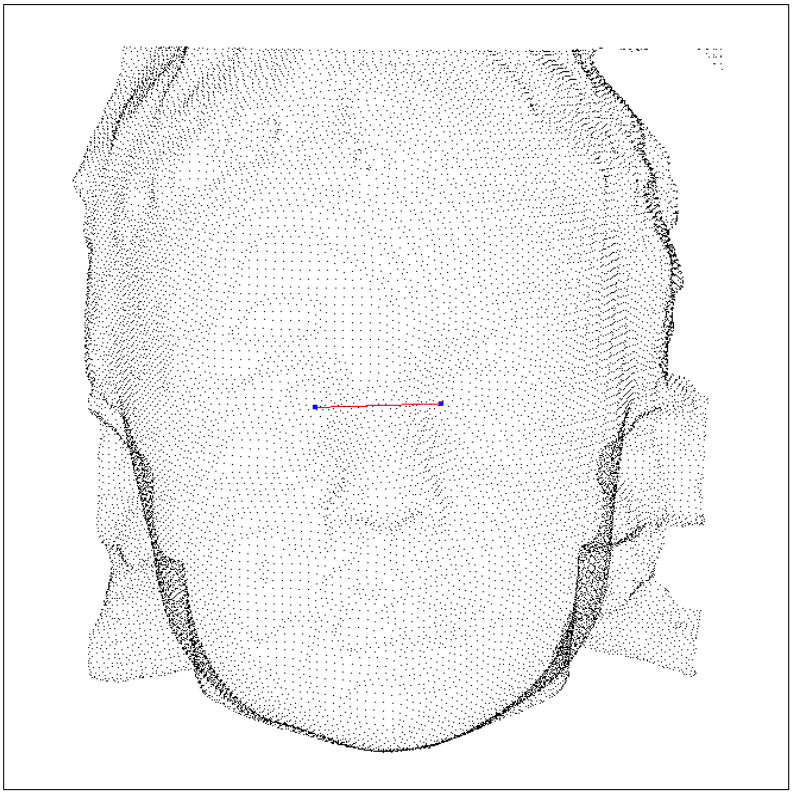

Figure 4.17: Directional vector between the eyes.

The effect of using this direction instead of the estimated minimum principal curvature direction is that a ridge point is guaranteed to be identified at regularly vertically spaced points, which should prevent:

(a) the same ridge point being identified more than once

(b) errors in principal direction estimation leading to ridge points not relevant to the required curve, for example points outwith the area of interest.

It could be argued, however, that by using a direction vector which is arbitrary, to the extent that it is independent of surface curvature, the extracted points are not strictly ridge points. This is certainly true, but the rationale is that we hope to identify a series of points corresponding to (in the illustration) the midline curve, each of which is a local maximum curvature point, though not strictly a ridge point. With regards to the direction, it is desired to have a series of points describing the curve as it progresses down the face, and so it seems natural to allow the direction of consideration to cross the face horizontally, and hence to pick out the ridge points which cross the midline plane roughly orthogonally. It seems acceptable, then, to make use of this horizontal (when the face is viewed from the front) vector for the midline curve, and to use a similar vertical vector (defined by the top and bottom landmarks of the nose) for the facial curves which cross the face roughly from left to right. The result of using this horizontal eye-vector is the series of ridge points shown in Figure 4.18, which seems to identify similar characteristics to the previous method, although this series of points should be more appropriate for smoothing, as the points will exist for every 'horizontal' line crossing the face, and will be, to a certain extent, more evenly spaced.

The main characteristics which are obvious from these points are the existence of ridge points to the side of the bridge of the nose, and around the sides of the philtrum area. This is entirely to be expected. The sides if the philtrum are raised and the method's validity could justifiably be questioned if these points were not to be identified. The ridge points identified at the sides of the nose are

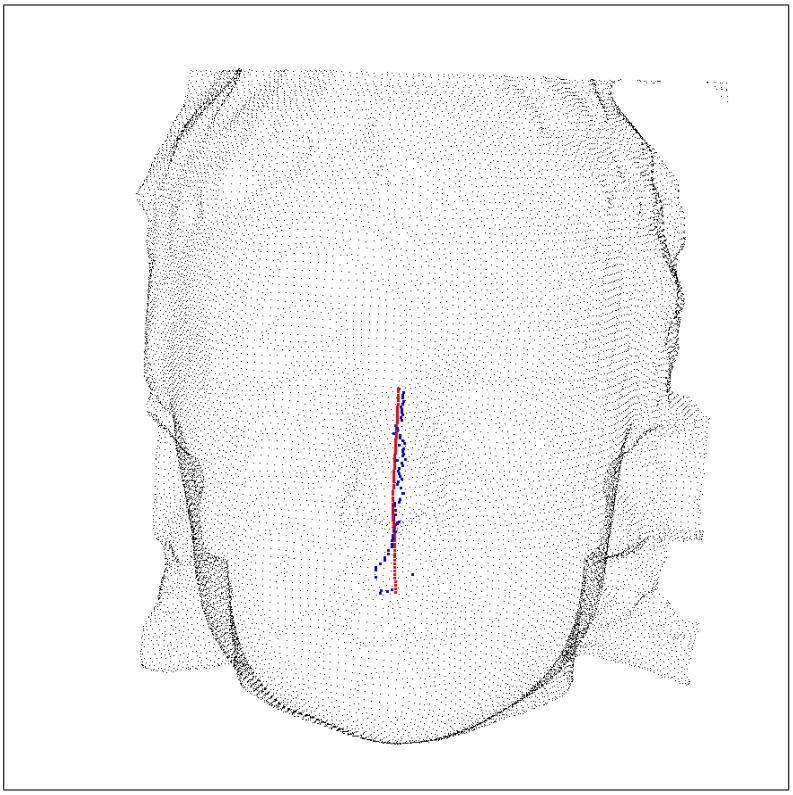

Figure 4.18: Ridge points identified from starting points of the planar midline curve using the direction based on the vector between the eyes shown in blue, with original planar curve points shown in red.

less intuitively correct, but it is clear on facial palpation that in fact the top half of the nasal bridge is actually fairly flat down the midline, with sharper changes in direction a few millimetres to either side. These are interesting facial features, which we shall return to discuss in Section 4.6.

Meanwhile, the current aim is to identify a midline curve based on the surface curvature. Some consideration of the absolute value of curvature for particular series of mesh points leads to the conclusion that in some cases the predicted cubic

smoothing spline has, for a portion of the points at least, constant curvature, such as the points and smoothing spline shown in Figure 4.19 along with the corresponding $\boldsymbol{\kappa 2}[\boldsymbol{\alpha}](t)$ value at each point.

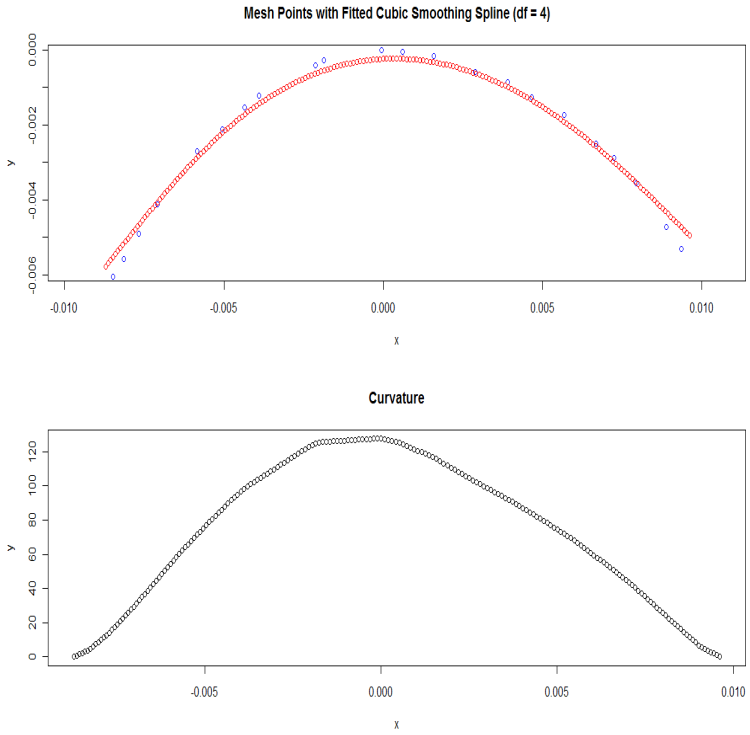

Figure 4.19: Curvature shown for a fitted cubic smoothing spline of degree 4. Notice the 'flat' top of the curvature points.

Notice the series of adjacent points which have almost constant curvature. This could be a problem in identifying the point of maximum curvature since tiny fluctuations in curvature value could lead to the identified ridge points being significantly different.

Therefore, to allow for multiple points with high, relevant curvature absolute

values, the method allows the identification of any point with sufficiently high curvature, defined in this case to be greater than 95% of the maximum absolute value of curvature, and then, for each planar starting point, to identify which point of high curvature lies closest to the original planar curve point. Effectively what this achieves, for identifying the corresponding curve to the planar curve, is to accept the already defined planar points providing their curvature was sufficiently high. This prevents small fluctuations in curvature producing estimates of ridge points which are wildly inaccurate. The results of this method are shown in the first image in Figure 4.20.

It is obvious that the points are reasonably irregular, and it does not seem sensible to stop at this point, and consider these (somewhat scattered) points as our extracted curve. A natural progression is to smooth the points so that the curve itself is smoother, and the points making up the curve are regularized, which will be more appropriate for analysis purposes.

A smooth curve was defined by a series of predicted points from a cubic smoothing spline, obtained using the methodology described at the start of Chapter 3 , subject to the constraint that the points corresponding to the facial landmarks on the midline curve must be fixed points on the smoothed curve. The smoothed curve is illustrated in the second two images in Figure 4.20.

A significant number of refining steps were necessary for this method of extracting a curvature based midline curve to perform sufficiently well, with the exception of the area around the philtrum, where the ridges are so sharply defined that the points around the midline curve have low curvature in comparison. The extension of this method to the other curves of the face would be a natural follow-on from this work, and would involve consideration of the particular area

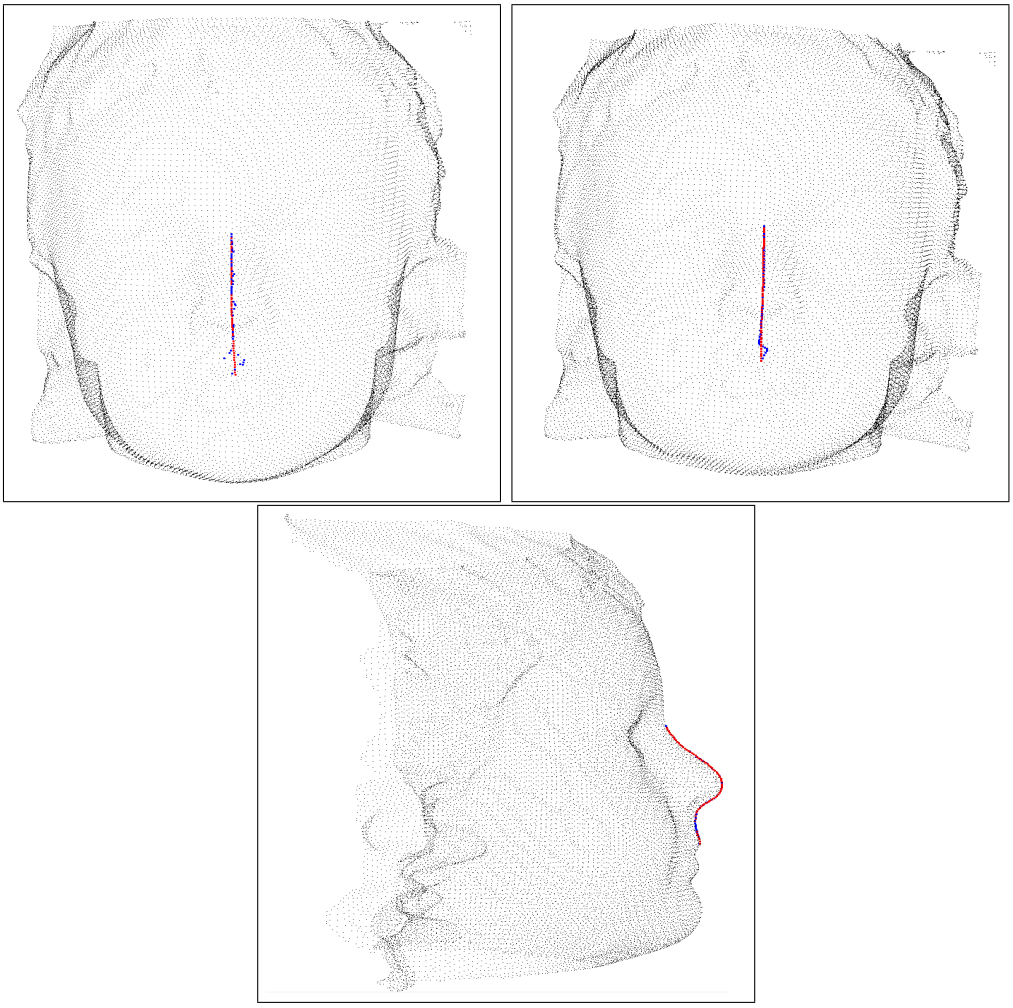

Figure 4.20: Extracted then smoothed midline curve, with ridge points extracted using the '> 95%' and 'horizontal direction' method, shown in blue, and the original planar curve point in red.

of the face of interest. The distinct characteristics of the different facial curves mean that this work is outwith the scope of this thesis, however the findings from this chapter should prove a solid foundation from which to begin. Of particular interest would be the principal components analysis of these curves, and comparison with the results obtained from the landmark-based analysis in Chapter 2 and the planar curve-based analysis in Chapter 3.

4.6 Facial Characteristics Identified during development of the Surface Curvature method

One of the aims of developing the surface curvature-based methods discussed earlier in this chapter was to allow a degree of flexibility in facial curve extraction by removing the constraint that the curve must lie on a single plane. However, it was also the intention to look carefully at the results of the intermediate steps to identify any interesting facial characteristics which are not immediately obvious from a landmark or standard curve based analysis. There were three main areas of interest which emerged:

- The identification, in the vast majority of subjects, of an apparent 'double ridge' tracking either side of the upper section of the bridge of the nose.
- The area of the philtrum, which is not considered as part of the standard set of facial curves, but is a distinctive area of the face, and in particular tends to be an area often affected in the faces of subjects who have had corrective surgery to repair a cleft lip and/or palate.
- The images in Figures 4.5, 4.6 and 4.7 provide extremely interesting views of the face, and certainly provide areas for discussion.

4.6.1 The double ridge

Looking back at Figure 4.19, the most interesting characteristic is that the graph of the curvature of the nose has a 'flat' top. A similar pattern was observed for, broadly, the upper half of the bridge of the nose. The implication from this is that the central area of the nose (as we move across the upper section of the

nose from one side to the other) has a largely constant, high degree of curvature when compared to the areas further from the bridge. Palpation of the nose confirms that this seems a reasonable finding, and that rather than one central ridge running down the bridge of the nose, in fact there seem to be two ridges running from the top to around halfway down the nose. This goes some way to explaining why, when extracting the midline nasal curve, it was not possible to simply identify ridge points. In fact, doing so provides a result similar to the blue points extracted in Figure 4.16, where many of the points identified as ridge points lie to the side of the central nasal bridge.

Additionally, if we consider the graph of curvature shown in Figure 4.12, where the spline used has a relatively high 8 degrees of freedom, the ridge characteristics can be seen even more clearly, with obvious peaks in curvature appearing either side of the midline. It seems that the smaller peaks in curvature, which occur to varying degrees for different areas of the nose, and in different subjects, are representative of other normal features, for example the roundness of the outside of the nostrils.

It is encouraging that the surface curvature analysis is detecting these characteristics, and although in the context of studies into cleft lip and/or palate, these features are not of particular interest, alternative areas of study may find the identification of these ridges useful.

4.6.2 The philtrum

Of more interest in the context of cleft studies is the nature of the area around the lips, particularly the area between the upper lip and the nose, where deformities are most apparent. The philtrum is the small flat area below the nose which

is bounded by two vertical ridges. Since cleft lip involves the deformity of the upper lip, the identification of the philtral ridges is certainly of interest in this context.

Figures 4.16, 4.18 and 4.20 all show that ridge points are being identified on the ridges of the philtrum because these (very distinct in control faces) ridges lie so close to the midline area under investigation.

4.6.3 Observations from the surface curvature images

This section will consider the illustrative images in Figures 4.5, 4.6 and 4.7, i.e. those displaying the maximum principal curvature, minimum principal curvature and Gaussian curvature for example faces. It is encouraging that the maximum and minimum principal curvature characteristics are very similar between subjects, as we would expect for surfaces as distinct as the human face. Features of note are as follows:

- The colours representing the principal curvature at each point are on the same scale for maximum and minimum principal curvatures, and so it is immediately apparent that the values of maximum principal curvature are generally in the upper half of the numerical scale, while the values of minimum principal curvature are generally in the lower half. This is what we would expect, and gives encouragement that the algorithm is working as intended. It should be remembered that a negative principal curvature value indicates bending away from the tangent plane at the point towards the face, while a positive principal curvature value indicates bending away from the face.

- The tip of the nose area can be seen to have distinct curvature characteristics. The tip of the nose has relatively low (negative) maximum curvature, as illustrated by the greens, oranges and reds in Figure 4.5, and very low (more negative) minimum curvature, as illustrated by the red area in Figure 4.6. We had previously conjectured that the tip of the nose may be close to being an umbilic, ie $k_{max} = k_{min}$, and this may still be a valid theory for those points where both maximum and minimum principal curvature are red. In general, though, for this area, we have $k_{min} < k_{max} < 0$, and this combination manifests itself in Figure 4.7 with the red elliptic concave Gaussian curvature area around the nose tip. An interpretation of this is that the tip of the nose is an area completely bending towards the face from the tangent plane at these points, or equivalently, that the area "sticks out" of the face. By considering other elliptic concave areas in Figure 4.7, we can see that the cheeks, forehead and areas of the lips and chin have similar characteristics.

- The eye sockets are also distinct. Figure 4.5 shows high values of maximum principal curvature in this area and Figure 4.6 shows that the minimum principal curvature values in this area are high, and at some points, positive. This corresponds to $k_{max} > k_{min} > 0$, i.e. the Gaussian curvature is elliptic convex, shown by the green areas in Figure 4.7.

- The majority of the facial surface has hyperbolic Gaussian curvature, i.e $k_{max} > 0 > k_{min}$. This corresponds to "saddle-shaped" areas of the face, where the maximum and minimum principal curvatures are bending in different directions away from the tangent plane to the point.

Chapter 5

Discussion

This thesis can be considered as two areas. The first area is the initial landmark-based analysis of the faces of 88 5-year-olds, acting as control cases for studies into cleft lip or cleft lip and palate. This analysis was considered in Chapter 2. The second area, consisting of the majority of the thesis, is an investigation into methods for extracting, or identifying, the curves of the face. This chapter will discuss the findings and conclusions of these two areas in sections 5.1 and 5.2 respectively.

5.1 Landmark-based Analysis

The initial analysis of the facial data, in Chapter 2, was based on a set of facial landmarks which had been identified by a trained operator. A source of human error was therefore introduced, and it was desired to quantify this error. The analysis based on landmarks was carried out using Procrustes Analysis, which allows the effects of location, scale and rotation to be removed from the objects so

that any remaining differences are due to shape. It was found that the individual operator who identified the landmarks for the cohort of interest tended to be good at repeated identification of any landmark.

There were some concerns, however, that arose from the result that the shape of the landmark configuration in the data of interest appeared to differ from the landmark configuration of a previous cohort of 5-year-old control faces. These faces had been marked up by a different operator, so a possible cause for the differences could be differences in the operators' interpretation of landmark definitions. An alternative explanation would be that the two cohorts came from different populations, but there is no obvious explanation as to why this would be the case. The multivariate analysis of variance used for this comparison should be interpreted with caution, however, due to the small sample size, particularly in the context of data with such high dimensions.

The landmark-based analysis then went on to consider whether any sexual dimorphism was present in the data, that is, whether any difference existed between the faces of the males and females of the cohort. Size was initially considered, and although the mean size of the male and female faces was similar, the spread of the sizes extended much higher for males than for females, and indeed a significant difference was found to exist between the size of the male and female faces. Once this effect of size was removed, it was then of interest to look for any difference in the shape of the landmark configurations of males and females. Once the effect of size had been removed, no significant shape difference was observed, which is in agreement with previous literature considering children of this age. A difference in shape between males and females would be expected in older children.

5.1.1 Areas for further study

It would be of interest to conclude why differences have been observed between the cohort of interest and a previous 5-year-old control cohort. In order to establish if these differences are due to the operators' identification of the landmarks, one of the following studies would be informative:

1. A single operator could mark up both cohorts.

2. Both operators could mark up a single cohort.

In either case, the multivariate analysis of variance could be repeated. If differences were still observed between the cohorts in 1, then we could perhaps conclude that these differences are not caused by operator interpretation of landmark definitions, and further investigation into the populations would be useful. If differences were observed between the two marked-up versions of a single cohort in 2, then we could tentatively conclude that the operators identifying the landmarks have different interpretations of the landmarks definitions.

In either case, the small sample size is a concern, and for this analysis it would be preferable to have a larger number of faces for analysis. However, collecting the data and marking up the landmarks is a very time-intensive process, and so this may be an unrealistic option.

With regard to the sexual dimorphism analysis, it would certainly be of interest to follow this cohort as they grow older, and investigate at what point male and female faces begin to show significant differences.

The natural progression for all of the landmark-based analysis is to move on to use the curves of the face for a similar analysis. The bulk of this thesis

involves methods for the extraction of these curves with the intention that the curves could be analysed in the future.

5.2 Curve Extraction

Methods for identifying the curves of the face are discussed in Chapters 3 and 4. Chapter 3 proposes a method for extracting curves based on the assumption that each facial curve for identification, or each of the component sections for those curves that are complex enough to require splitting, lies on a single plane in three-dimensional space. Chapter 4 relaxes this assumption and discusses methods for identifying the facial curves in question using the surface curvature characteristics of the face.

The method for extracting facial curves using planes involved two steps. Firstly the plane upon which the curve (or section of curve) was to be constrained to lie was defined. Simple algebra then allowed the identification of a series of irregular points which comprised all intersections of the facial mesh and the plane. A cubic smoothing spline was then fitted to these points to allow them to be regularized.

Two complementary methods for defining the plane were employed, since unfortunately neither method was successful for all facial curves (or curve sections). The first method defined the (unique) plane containing three named landmarks which were all to lie on the plane. This method did not work for the curve extraction algorithm when the defined plane was close to being parallel to the face, as it then intersected many areas of the face, including areas very near to the curve of interest, which caused problems. However, this method did work well

for a number of the facial curves.

The second method, designed to prevent the plane cutting the face at a problematic angle, involved the use of two landmarks to lie on the plane, with one landmark forced to lie orthogonally to the plane. This again produced a unique plane, but worked well for those curves which had not been defined well using the first method.

The planar method of curve extraction worked well, although there are a couple of areas which would benefit from future refinement. These are discussed in Section 5.2.1.

The method of extracting curves using the surface curvature characteristics of the face is discussed in Chapter 4. The main theory underlying this chapter involves the identification of the principal curvatures and principal directions at each point on the facial mesh. Considering the tangent plane to any point, $\mathbf{p}$, on a surface $\mathcal{M}$, we can consider any tangent vector $\mathbf{v_p}$, lying on the tangent plane, passing through $\mathbf{p}$. For each tangent vector, the normal curvature $k(\mathbf{v_p})$ is defined, which quantifies the extent of bending of the surface $\mathcal{M}$ in the direction $k(\mathbf{v_p})$. There are therefore many values of normal curvature at $\mathbf{p}$, in every possible direction within the tangent plane. The directions which correspond to the maximum and minimum normal curvatures are the principal curvatures, k_{max} and k_{min}, and their corresponding tangent vectors are the principal directions. Another interesting measure of curvature, the Gaussian curvature, is equal to $k_{max}k_{min}$, and we can define ridge lines as curves on a surface along which the surface bends sharply. Ridge lines consist of a series of *ridge points*. Ridge points can also be defined in terms of principal curvatures and principal directions. It can be seen that principal curvatures and directions form a strong basis

for investigating many aspects of the curvature of a surface.

To estimate the principal curvatures at any point, a cubic surface, known as a *Monge patch* is fitted to the local facial mesh.

The theory described above was implemented and illustrated with the facial midline curve, as follows: The maximum and minimum principal curvatures at each point were calculated, as well as the corresponding Gaussian curvature, and this allowed visualisation of the face and consideration of the obvious features. The ridge points were then identified by looking for the points of greatest curvature along the minimum principal curvature direction. It was of interest to establish whether any ridges corresponded to the facial curve of interest. Using a combination of cubic smoothing splines with different degrees of freedom, it was possible to identify a number of ridge points in the region of the midline curve. However, these were found to form a pattern which was unlike the midline curve of interest.

A refinement of the method was implemented, which used a horizontal (i.e. crossing the face from left to right) direction rather than the minimum curvature direction to define a series of points which were points of local maximum curvature of the face when moving from left to right. These points were a slight improvement on the ridge points in terms of identifying a midline curve, since they were broadly evenly spaced and appeared slightly less irregular than the ridge points. However they were still not an appropriate set of points for defining the midline curve, particularly because of the variation in the area around the philtrum, where the philtral edges were identified.

It became apparent that one of the reasons for the midline point being identified was that particular plane curve had areas of almost constant minimum

curvature, extending across a umber of points, in most cases spanning the bridge of the nose. When this occurred, tiny fluctuations in the calculation of curvature may completely change the location of the identified ridge point.

A further refinement was therefore introduced, which was to consider any point on a plane curve with *sufficiently high* curvature (i.e. greater than 95% of the maximum curvature) to be a potential ridge point, and then of these high curvature points, for each horizontal crossing, to choose the point which lay closest to the original midline curve. This has the result of accepting the current planar curve points, provided their curvature was sufficiently high.

This method proved successful in identifying a midline curve, with the exception of the area around the philtrum, since the philtral edges are points of high curvature, and are so sharply defined that the points near the midline curve in this area have relatively low curvature in comparison. The development of this method will allow further steps to be taken in this area, and these are outlined in Section 5.2.1.

Chapter 4 also provided an opportunity to consider the curvature of the face as an interesting topic in in its own right.

5.2.1 Areas for further study

The planar method for extracting the facial curves could benefit from further work, in particular:

- It would be preferable to have only a single method for defining the plane on which the curve should lie. This would simplify the algorithm for the extraction.

- The algorithm, written in the form of R code, would benefit from modifications to reduce computational time.

In the context of the surface curvature method for extracting curves, a number of areas for future study are proposed:

- Further refinement of the method for extracting the midline curve to allow correct identification of the area around the philtrum, perhaps by looking for blue ridges (valleys) rather than red ridges in this area.

- An extension of the surface curvature method to the other facial curves. This will involve consideration of each curve, or section, in turn to choose appropriate directional vectors and local areas for consideration.

- Development of the R code used for the extraction algorithm to ensure computational efficiency.

- The analysis of a set of extracted curves would be the natural conclusion to this study. Barry (2008) presents a longitudinal method for analysing shape data in a facial context, which could be implemented using curves extracted from the cohort at different ages.

Bibliography

Abrahams, P. H., Hutchings, R. T. & Marks Jr, S. C. (1998), *McMinn's Colour Atlas of Human Anatomy*, fourth edn, Mosby.

Ayoub, A. F., Garrahy, A., Hood, C. A., White, J. E., Bock, M., Siebert, J. P., Spencer, R. & Ray, A. (2003), 'Validation of a Vision-Based, three-dimensional facial imaging system', *Cleft Palate-Craniofacial Journal* **40**, 523–529.

Barry, S. (2008), Longitudinal Analysis of Three-dimensional Facial Shape Data, PhD thesis, University of Glasgow.

Bock, M. & Bowman, A. W. (2006), 'On the measurement and analysis of asymmetry with applications to facial modelling', *Applied Statistics* **55**(Part 1), 77–91.

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

Dryden, I. L. & Mardia, K. V. (1998), *Statistical Shape Analysis*, John Wiley and Sons, Chichester, UK.

Farkas, L. G. (1994), *Anthropometry of the Head and Face*, second edn, New York: Raven Press.

Ferrario, V. F. & Sforza, C. (1997), 'Size and shape of soft-tissue facial profile: Effects of age, gender, and skeletal class', *Cleft Palate Craniofacial Journal* **34**(6), 498–504.

Ferrario, V. F., Sforza, C., Schmitz, J. H., Miani, A. J. & Taroni, G. (1995), 'Fourier analysis of human soft tissue facial shape: sex differences in normal adults', *Journal of Anatomy* **187**, 593–602.

Goldfeather, J. & Interrante, V. (2004), 'A Novel Cubic-Order Algorithm for Approximating Principal Direction Vectors', *ACM Transactions on Graphics* **23**(1), 45–63.

Gray, A., Abbena, E. & Salamon, S. (2006), *Modern Differential Geometry of Curves and Surfaces with Mathematica*, Chapman and Hall/CRC, Boca Raton.

Hallinan, P. L., Gordon, G. G., Yuille, A. L., Giblin, P. & Mumford, D. (1999), *Two- and Three- Dimensional Patterns of the Face*, A K Peters, Natick, Massachusetts.

Hastie, T. J. & Tibshirani, R. J. (1996), *Generalized Additive Models*, Chapman and Hall, London.

Kopka, H. & Daly, P. W. (2004), *Guide to LaTeX*, fourth edn, Addison Wesley.

Kühnel, W. & Hunt, B. (2006), *Differential Geometry. Curves - Surfaces - Manifolds*, American Mathematical Society, United States of America.

Ohtake, Y., Belyaev, A. & Seidel, H.-P. (2004), 'Ridge-Valley Lines on Meshes via Implicit Surface Fitting', *ACM Transactions on Graphics* **23**(3), 609–612.

Page, D. L., Koschan, A. F., Abidi, M. A. & Overholt, J. L. (2006), 'Ridge-Valley Path Planning for 3D Terrains', *Proceedings of the 2006 IEEE International Conference on Robotics and Automation* **1–10**, 119–124.

Prahl-Andersen, B., Ligthelm-Bakker, A. S. W. M. R., Wattel, E. & Nanda, R. (1995), 'Adolescent growth changes in soft tissue profile', *American Journal of Orthodontics and Dentofacial Orthopedics* **107**(5), 476–483.

Venables, W. N. & Ripley, B. D. (1997), *Modern Applied Statistics with S-Plus*, Springer-Verlag, New York.

Venables, W. N. & Ripley, B. D. (2000), *S Programming*, Springer, New York.

White, J. E., Ayoub, A. F., Hosey, M. T., Bock, M., Bowman, A. W., Bowman, J., Siebert, J. P. & Ray, A. (2004), 'Three-Dimensional Facial Characteristics of Caucasian Infants Without Cleft and Correlation With Body Measurements', *Cleft Palate-Craniofacial Journal* **41**, 593–602.

Yamada, T., Mori, Y., Minami, K. & Sakuda, M. (2002), 'Three-dimensional analysis of facial morphology in normal Japanese children as control data for cleft surgery', *Cleft Palate Craniofacial Journal* **39**, 517–526.