Longitudinal Study of Facial Growth

Anna Price

A Dissertation Submitted to the
University of Glasgow
for the degree of
Master of Science

School of Mathematics & Statistics

October 2013

Abstract

The statistical analysis of shape is an area with wide applications. The field of facial shape analysis is particularly interesting because the human face has such a complex shape. This thesis investigates the change in the facial shape of a cohort of children, whose faces were photographed 6 times between the ages of 3 months and 12 years providing a normative control group for potential future comparisons.

Chapter 1 provides an introduction to the data, detailing the background and motivation for the longitudinal study. The aims of this thesis are outlines.

Chapter 2 describes the study design and data in detail. The process of data collection, both in the original study and the 12-year follow-up, is explained and the resulting data is explained. Issues that arose with data collection and the resulting data are introduced and discussed.

In Chapter 3 the theory of various statistical shape analysis methods is described in detail, including Procrustes Analysis and Principal Component Analysis.

Chapter 4 develops and discusses methods of automatic curve identification, as well as providing the theory of the shape index and methods of fitting curves such as p-splines and principal curves.

Chapter 5 details the longitudinal analysis of the control data. Statistical methods such as Procrustes Matching and Linear Mixed Effect models are implemented to discover change in facial shape over time.

In Chapter 6 the findings of the thesis are discussed, as well as the limitations of the data and the statistical methods used. Further work which

would be of interest in the future is proposed.

Acknowledgements

I would like to thank my supervisor Prof. Adrian Bowman for his invaluable guidance and support over the past year. I would also like to thank all of the children and parents who returned to participate in the follow-up image capturing sessions with such enthusiasm. Without them this thesis would not have been possible.

I am very grateful to the Engineering and Physical Sciences Research Council for funding my research.

Finally, I thank my fellow post-grads for providing distractions when they were needed most (and when they weren't) and keeping me sane for the year.

Declaration

I have prepared this thesis myself; no section of it has been submitted previously as part of any application for a degree. I carried out the work reported in it, except where otherwise stated.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Introduction

The human face contains a huge amount of information and can tell you a lot about a person, for example it may indicate their gender, age or even ethnic background. However, it can be difficult to identify what about a person's facial structure tells us this kind of information. Although the human face is an instantly recognisable shape, it is also an extremely complex one. Describing the surface of a three-dimensional shape is a difficult task, but achieving an accurate description can assist work in many different applications, such as in medical research and practise, physical anthropology and the development of automatic facial recognition systems. In the medical field, accurate information about facial shape is particularly advantageous for facial surgeries.

The Face 3D Project is a research group involving several institutions and is funded by the Wellcome Trust. The consortium consists of researchers from the University of Glasgow (Statistics, Computer Science and Dental Science), the Royal College of Surgeons in Ireland, Dublin City University, the University of Limerick and the Institute of Technology in Tralee. The Face 3D Project aims to develop new methods of extracting and analysing facial shape information from three-dimensional images. These methods are

then applied to data in two application areas. The first is the characterisation of the biological process underlying schizophrenia. The second area of research focus is the quantative assessment of the outcome of facial surgeries, specifically surgery of the jaw - orthognathic surgery.

1.2 Motivation for the Study

In the field of statistical facial shape analysis, much research has focused on two-dimensional data (Bock and Bowman (2006), Kent *et al.* (2001), Morris *et al.* (2000) and Morris *et al.* (1999)). However, there has been considerably less work carried out on three-dimensional facial shape analysis. As discussed above, an understanding of facial shape and growth can aid in the planning and evaluation of facial surgery for patients with craniofacial anomalies. The extent of the facial deformity can be assessed if a baseline is available for comparison. Therefore it is of interest to characterise normal facial features in three-dimensions, and to understand how facial shape changes as a person grows from infant to adult.

One example of a craniofacial abnormality for which this kind of information is vital is cleft lip and palate which, according to the NHS, affects one in every seven hundred babies born in the UK (NHS, 2012). Cleft lip and palate is characterised by a gap or split in either the upper lip or the roof of the mouth (the palate), or sometimes both. The aim of cleft lip and palate surgery in children is to improve the appearance and function of the nose and mouth area and to achieve a facial shape that is similar to that of normal subjects. In this case it is essential to establish a standard model from normative control data that allows for comparison of facial shape between normal children and children with cleft lip and palate. It is also important to understand how the facial shape changes over time in children with cleft lip and palate pre- and post-surgery. Again, it is a requirement to have a baseline for comparison of facial development. There have already

been some studies using normative data in Japanese children. Yamada *et al.* (2001) analysed the facial forms of normal Japanese children aged between 4 months and 3.5 years. Kau *et al.* (2006) constructed a model of the average facial shape of Japanese children with a mean age of 11.8 years. However, there are no studies that have similarly described the facial shape and growth of normal Caucasian children.

1.3 Collecting Facial Data from Infants

There are some problems that naturally arise when collecting data from infants. A child's age, the level of cooperation given by each child and the length of the examination session can all affect the quality and accuracy of the data. Movement is a key issue when collecting facial data as it is unlikely that an infant will be able to hold one specific position for longer than a few seconds. Previous methods of data collection have included direct anthropometry, two-dimensional photographs, radiographs and facial impression models, but none of these is able to adequately deal with the aforementioned issues while simultaneously providing high quality data. Furthermore, some of these techniques require the child to undergo general anaesthesia or expose the child to radiation, both of which place the child at unnecessary risk. One method which resolves all of these problems is to use three-dimensional scanning equipment. This was the method employed by Yamada *et al.* (2001) where a liquid-crystal range finder scanned the face for 1 second, producing three-dimensional coordinates and colour information. Similarly, a recently developed computerized stereophotogrammetry camera system (the Di3D system and software) allows the capture of monochrome and colour images in just 50 milliseconds (Ayoub *et al.* (2002)). Firstly, the images are captured exceptionally fast and hence the child is not required to be still in one position for longer than a few seconds. Secondly, the system is flexible as it allows multiple images to be captured in one session - this is especially

useful if the positioning or facial expression of the child is not correct in previous pictures. Finally, the method is non-invasive and poses no significant risk to the child.

1.4 Areas for Investigation

This thesis focuses on three main areas of study. The first is the process of data collection of three-dimensional facial images from a cohort of 12-year old children. These subjects first participated in a longitudinal study of facial shape in 2000, and the design of the original study will also be described. The second area of focus is the exploration of methods of extracting information from a three-dimensional facial image. Specifically, the aim is to produce three-dimensional curves which accurately capture key features of the face. The final section of this thesis uses one of the methods explored in order to analyse the data collected in the longitudinal study. Linear mixed models will be fitted in order to identify what contributes to the change in facial shape over time and to explore whether there are differences between the sexes.

1.5 Summary

The applications of statistical shape analysis are greatly varied, and one such field is the study of facial shape. In order to improve facial surgery it is necessary to accurately describe and analyse a facial surface. The Face 3D Project provides an opportunity to study the development of facial shape over time in a cohort of children using innovative camera technology. This thesis describes the study design and data available, explores methods of obtaining three-dimensional curves and finally uses these methods in order to investigate the data.

Chapter 2

Study Design and Data

2.1 Design of the Longitudinal Study

In 2000, a longitudinal study of facial growth in healthy infants began at the Glasgow Dental Hospital & School, using the innovative Di3D camera system discussed in Section 1.3. The aim was to establish normal facial shape for children at 3 months, 6 months, 1 year and 2 years of age for comparison with a sample of children with unilateral cleft lips (UCL) or unilateral left lip and palate (UCLP). Some results published from the initial data collection at 3 months showed a significant difference in facial dimensions, but little difference in facial shape between the sexes (White *et al.* (2003)). The study organisers chose to follow-up this cohort of children at 5 years of age. In 2012 it was again decided to follow-up the control group, who were then 12 years of age, in order to enlarge the longitudinal data set and eventually carry out longitudinal analysis from 3-months to 12 years.

In the original data collection process, it was estimated that a control group of 60 infants would be necessary in order to achieve three times the number of controls as cleft cases. Additionally, a high dropout rate was anticipated and so 100 infants in total were recruited to the control group. The subject requirements were that the infants had to be Caucasian, delivered at full term with no perinatal complications, no craniofacial abnor-

malities and with a birth weight above the third percentile and below the 97th percentile, compared to growth and development records of Gairdner and Pearson (Gairdner and Pearson (1988a); Gairdner and Pearson (1988b)). These requirements ensured that the study group was homogenous and able to provide normative data.

Recruitment to the study was carried out by placing leaflets and posters in local maternity wards, health centres, general medical practitioner surgeries, the Glasgow Dental Hospital & School and the Royal Hospital for Sick Children, Glasgow. Furthermore, parents were personally invited to participate in the study by the study organisers at maternity wards and immunisation clinics. Parents were asked to complete a consent form and a questionnaire which gathered information on contact details, place of birth, infant's birth weight, date of birth and whether or not the infant was breast feeding. Family social status was assessed using the Carstairs Index of Deprivation. Details for a secondary contact, such as the grandparents or an aunt or uncle of the child, were collected at the initial meeting and this proved very valuable when following up subjects at later time points.

Of the 100 subjects who were initially recruited, 42 changed address at least once during the study, with several families moving away from the Glasgow area or even abroad. The resulting data consists of three-dimensional images of the infants at 3 months, 6 months, 1 year and 2 years of age. These were captured using a three-dimensional computerised imaging system situated at the Glasgow Dental Hospital & School and at the Royal Hospital for Sick Children. In 2005 the study team decided to follow-up the children, then aged 5 years. At this time point 88 subjects were successfully followed-up.

2.2 Design of the 12-year Follow-up

It was decided in 2012 to follow-up as many subjects from the initial study as possible. A database of contact details was compiled from the

original forms filled in by parents at baseline, additional forms filled in by parents at the 5-year visit and a small database of email addresses for some of the parents. There were details on a total of 62 of the 88 subjects whose images had been captured at the 5-year time point. However, due to missing parents' names, three subjects were excluded from the list. This left a group of 59 potential subjects for follow-up at 12 years and it was considered that a response rate of two thirds of this group, or approximately 40 of the subjects, would be adequate for future longitudinal analysis.

Ethical approval was granted for the follow-up study by the Ethics Committee of the College of Science & Engineering in the University of Glasgow. The key ethical issue of the follow-up study was that the children had to be accompanied by a parent or guardian to the imaging session. This was made clear in the initial invitation letter to parents or guardians. Additionally, the details of the study were fully explained in the invitation letter along with the intended use of the data that would be collected. Finally, a consent form for each child was required to be filled out by a parent or guardian.

The first form of contact made with the parents was a letter to the last-known postal address of the family containing detailed information on the study (see Appendix A). If the child and parents were happy to continue their participation in the study, parents were invited to directly contact the study organisers or to fill out a form with details of their preferred method of contact. Stamped addressed envelopes for the School of Mathematics and Statistics were included in order to encourage the response rate should parents choose the latter form of contact. This was recommended by the previous study organiser from the initial recruitment process and did prove to be a preferred form of contact for many parents. After a few weeks, further invitations were sent by email to subjects for whom no response had been received and for whom an email address was available. Unfortunately, this method did not result in many responses as a large proportion of the addresses were no longer valid and the emails were returned. In a further

attempt to contact subjects, letters were sent to the person nominated as a secondary contact at the initial baseline visit. This letter explained that we were trying to contact the family regarding a study of facial shape that they previously participated in and requested whether up-to-date contact details could be provided, should the person be happy to do this. Again, the response rate to these letters was very low, probably because the details were at least 7 years old and so many of the contacts may have moved address. Finally, a second invitation letter was sent to the parents who had still not replied, to prompt a response. A further issue in arranging follow-up sessions was that, even once parents had initially responded and shown interest in participation, it often took weeks, if not months, to fix a date and time. Initially, it was thought that all the images could be captured over a small number of days, but it became apparent that this was not practical given the time constraints and other commitments of both the parents and children. Therefore sessions were arranged as responses arrived and dates and times chosen according to the convenience of the parents and child.

Table 2.1: Summary of the Follow-up Process at the 12-year Timepoint

Potential Subjects for Follow-up	Requested Removal from Study	No Response	No Session Arranged	Successful Follow-up
59	1	15	4	**39**

Of the 59 potential subjects, one family requested to be removed from the study, leaving 58 children to attempt to follow-up. A number of families had moved again within the Glasgow area since the 5-year time point, and a couple had moved away from Glasgow. Fifteen parents did not respond to any attempts at contact and it eventually proved impossible to arrange a visit for 4 of the children, despite their parents initially expressing interest in participating in the study. This means that the final data set consists images captured from 39 subjects. A database of the 39 final subjects was created for future use, with the most up to date contact details that were given by

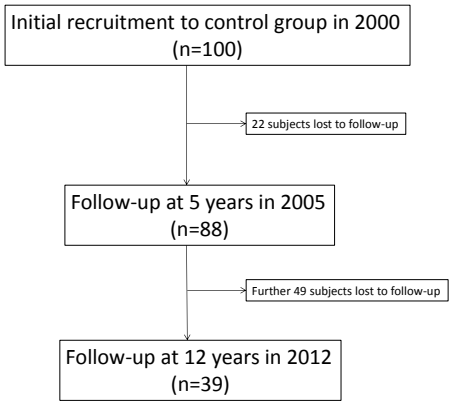

Figure 2.1: Flow diagram of subjects through the study

Figure 2.2: Di3D camera system

the parents. A diagram of the flow of participants through the study is shown in Figure 2.1.

2.3 Image Capturing

Images were captured in the Mathematics & Statistics building at the University of Glasgow where the Di3D camera system is situated. The system comprises four independent digital cameras attached around a static frame which sits on a large tripod. A picture of the camera system is shown in Figure 2.2. The four cameras capture images simultaneously and produce four standard digital images.

An example of the resulting jpeg images from one of the 12-year old

control children can be seen in Figure 2.3. The Di3D software builds the four images into a 3-dimensional image, as exemplified in Figure 2.3.

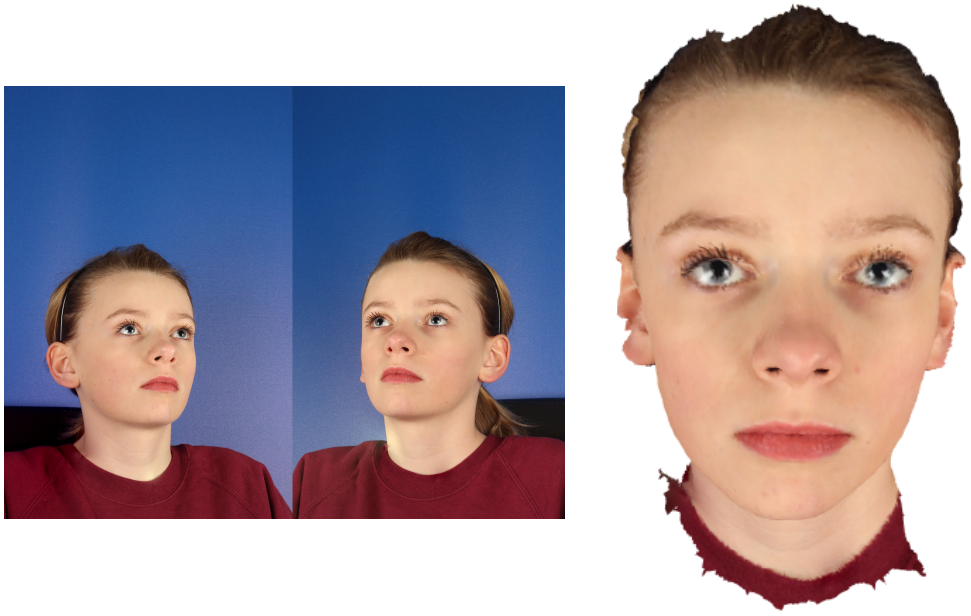

Figure 2.3: Digital stills produced using a Di3D camera system (left) and three-dimensional image (right)

The positioning of the camera was kept the same, or as close as possible, for every image captured. Before each session a set of images was taken of an object of known dimensions - in this case a rectangular board mounted on a small tripod. This allows the system to calibrate and this must be done before each session. A chair with adjustable height was placed directly in front of the camera.

At each session, a signed consent form was obtained from the parents or guardian present, which can be seen in Appendix B. The child's previous images, from 3 months, 6 months, 1 year, 3 years and 5 years of age, were displayed as interactive 3-dimensional images. The main reason for showing the previous images was to confirm correct matching of ID numbers, though the child was also encouraged to take advantage of the interactive nature of the images. This proved highly interesting and fun for both the child and the accompanying parents or guardian. If necessary, subjects were asked

to tie back hair with a hair band or clips and to remove jewellery such as earrings. The subject was seated in the chair in front of the camera and time was taken to achieve the correct positioning. This requires that all four cameras are centred on the middle of the subject's face. The subject was asked to tilt their child upwards slightly and look at the join between the opposite wall and ceiling. They were also asked to keep their lips closed, but relax their face completely as the clenching of the teeth or the holding of the mouth in an unnatural position leads to an inaccurate measurement of the true facial structure. In most cases, just one picture was required to obtain a satisfactory image. However, if necessary it was possible to take multiple pictures. At the end of the session, the possibility of future imaging sessions was mentioned to the parents. All the parents and children expressed interest in being contacted again in a few years, should the study organisers continue the longitudinal analysis. In the weeks following each session a copy of the image captured, along with copies of the five previous images, were sent to the family in the form of a short video as a memento to keep. The protocol document for the sessions can be found in the Appendix (C).

2.4 Data

In the original study, 100 infants were recruited at baseline and 88 children were followed-up to the 5-year time point. As discussed previously, there are 39 images captured at 12 years. Due to some issues with the data that was collected, 4 cases had to be excluded from analysis. The reasons for these exclusions is discussed in detail in Section 2.5.

Each image consists of approximately 19,000 three-dimensional coordinates which fully describe the surface of the face. Also available are groupings of three coordinates, or 'triples', which allow the face to be represented as a triangular mesh, where each triple gives the indices of the coordinates of the vertices of a triangle within the mesh. Furthermore, the colour to be

assigned to each point is available, which allows the face to be displayed in full colour which is useful when landmarking.

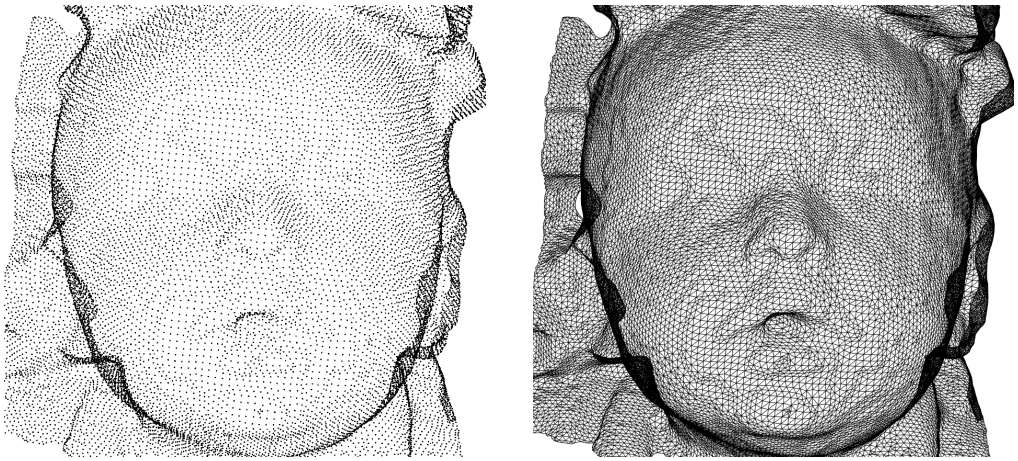

Figure 2.4: Facial image represented by three-dimensional coordinates (left) and a mesh of triples (right)

Figure 2.4 shows an example of a 3-dimensional image represented by coordinates and by a mesh of the triples. Each image that is captured by the camera system is loaded into a software tool called *Landmark*. This allows the facial image to be visualised, rotated and zoomed in three dimensions. Anatomical landmarks can be manually placed using this software. At the 12-year time point each image was allocated 30 anatomical landmarks, or 31 for the two cases in which the lips were slightly parted. Table 2.2 defines the landmarks that were allocated and Figure 2.5 shows an example image with the 30 anatomical landmarks superimposed.

2.5 Problems with the Data

One problem that arose with the 12-year old data was that, due to the age group of the children, a number had fixed braces. Clearly these cannot be removed for the image to be captured, but they have the potential to affect the facial shape. This was therefore noted in the database detailing information on the subjects. However, a further problem arose in two cases

Table 2.2: Anatomical Landmarks used for 12-year Follow-up Images

ID Number	Region of the Face	Type	Abbreviation	Name
1	nose	middle	pm	pronasale
2-3	nose	left/right	a	alare
4-5	nose	left/right	ac	alare crest
6	nose	middle	sn	subnasale
7-8	nose	left/right	nt	nostril top point, columella breakpoint
9	nose	middle	cc	columella constructed point
10-11	nose	left/right	nb	nostril base point
12	nose	middle	se	sellion
13	nose	middle	n	soft tissue nasion
14-15	eye	left/right	ex	exocanthion
16-17	eye	left/right	en	endocanthion
18-19	ear	left/right	t	tragion
20	lips	middle	ls	labiale superius
21-22	lips	left/right	cph	crista philtri
23-24	lips	left/right	ch	cheilion
25	lips	middle	sto	stomion
26	lips	middle	li	labiale inferius
27	chin	middle	sl	sublabiale
28	chin	middle	pg	soft tissue gnathion
29-30	ear	left/right	a	otobasion inferius

as the children had their lips parted in the image. The decision was taken to remove these cases from analysis, as the usual landmarks cannot be added to the lips. For these cases to be included in future analysis, an alternative method of landmarking would need to be employed. For example, two landmarks could be positioned in place of the stomion, one in the centre at the bottom of the upper lip and one in the centre at the top of the lower lip. Then a double "dummy" landmark would need to be added in all other cases, in order for the landmarking to be consistent and therefore comparisons to be made between subjects in the analysis.

Another problem that occurred in the 12-year old data was that two of the

Figure 2.5: Facial images with anatomical landmarks: front view (left) and side view (right)

images had small areas of distortion. This can happen as the Di3D software builds the four still images into a 3-dimensional image. As this occurs after the imaging session, it is not possible to take an additional alternative image. Unfortunately the areas of distortion occurred around the nose and mouth in both of the images. This meant that neither image was appropriate for analysis and also had to be excluded from the study.

One specific area of the face which can cause further problems is the ears. This can be due to hair around the ears which distorts the image, reducing the quality of the data and making it difficult to accurately place landmarks on the ears. Hair bands were provided in an attempt to pull hair away from the face and ears. However, in some cases it proved impossible to remove all hair from the picture. Distortion of the ears can also occur due to the positioning of the subject's ears in relation to the cameras. If the line of sight of the cameras is close to the angle of the face, the camera tends to capture the ears less accurately than the rest of the face. Due to this, analysis of the longitudinal data focused on the central features of the face, such as the

eyes, nose and mouth.

2.6 Summary

The data for this study consist of multiple 3-dimensional images collected from a cohort of 100 children who were recruited as infants in 2000. These images were captured at 6 different time points - 3 and 6 months, and 1, 2, 5 and 12 years. Due to drop outs the final cohort size at 12 years was 39. Each image provides a full description of the facial surface and allows for landmarks to be placed at key features of the face. Some problems that occurred with the data collected included fixed braces to the teeth and small distortions during the 3-dimensional image building process, resulting in a small number of unusable images.

Chapter 3

Statistical Shape Analysis

3.1 Measuring Shape

The aim of statistical shape analysis is to compare objects that have different shapes. This is because we are interested in exploring patterns of variability, or in finding the areas of greatest variability. In order to compare objects we require a way of measuring shape, a measure of distance between shapes and methods for statistical analysis.

Dryden and Mardia (1998), define shape as being "*all the geometric information about an object which remains after location, scale and rotational effects have been removed*". In other words, the shape is not affected by the three Euclidean similarity transformations (translation, rotation and reflection). In order to accurately represent an object's shape we allocate points, or landmarks, to specific locations on each object which correspond between and within populations. A landmark could represent a biologically meaningful position such as the corner of an eye (an anatomical landmark) or be assigned according to a mathematical requirement such as a point of high curvature (a mathematical landmark). The set of landmarks on a particular object is called a configuration and the $k \times m$ matrix of coordinates of the k landmarks in m dimensions is the configuration matrix X. We also define the pre-shape Z as being an object that has had information regarding location

and scale removed, i.e. rotational effects still apply. In the analysis of facial shape we take the number of dimensions m to be 3. Finally, the shape space is defined as the set of all possible shapes and is a non-linear space.

3.2 Measuring Distance

A commonly used measure of distance is Euclidean distance, which is simply the straight line distance between two points. However, the shape space is characterised by a curved surface and hence is non-Euclidean. Therefore, we require an alternative measure of distance: the Procrustes distance. If we have two configuration matrices X_1 and X_2 with pre-shapes Z_1 and Z_2, then the Procrustes distance is the closest great circle distance between Z_1 and Z_2 in the pre-shape space, with the minimization carried out over rotations.

3.3 Measuring Size

The most commonly used estimator of size in geometrical shape analysis is the centroid size, which gives the overall size of a shape in Procrustes analysis. The centroid size is the square root of the sum of squared distances of a set of landmarks from their centroid and is defined as follows:

$$S(X) = \sqrt{\sum_{i=1}^{k}\sum_{j=1}^{m}(X_{ij} - \bar{X}_j)^2}$$

where X_{ij} is the $(i,j)th$ entry of the configuration matrix X and $\bar{X}_j$ is the arithmetic mean of the jth dimension.

3.4 Procrustes Analysis

Now that shape and distance have been defined, we are interested in comparing the shapes of two or more objects. The landmarks that are assigned to each shape do not share the same origin and hence we cannot immediately

compare them to each other. In order that any 'real' shape difference can be identified (according to the definition of shape above), the configurations in a set must first be standardised so that they are of equal size, matched in rotation and centred at a common origin. The method of Procrustes matching optimally translates, rotates and scales the shapes to obtain similar placement. Ordinary Procrustes Analysis (OPA) is used to match two configurations to each other. Given two configurations X_1 and X_2 we have:

$$\min_{\beta,(R)} || X_1 - \beta X_2 R ||^2$$

where β is the scale parameter and R is the rotation matrix. Generalised Procrustes Analysis (GPA) is used to compare three or more shapes to an optimally determined mean shape. Given configurations $X_1 \cdots X_n$ we have:

$$\min_{\beta_i,(R_i)} || \beta_i X_i R_i - \bar{X} ||^2$$

where $\beta_1 \cdots \beta_n$ are the scale parameters, $R_1 \cdots R_n$ are the rotation matrices and $\bar{X}$ is the average of the set of configurations. GPA is an iterative process over all configurations in the set and $\bar{X}$ is given some constraint, such as $|| \bar{X} ||^2 = 1$. There are a number of variations on Procrustes matching. Full Procrustes matching involves matching configurations using all three similarity transformations, whereas partial Procrustes matching preserves the size of the objects by excluding scaling from the transformations. Both OPA and GPA can be extended to include reflection as well as the three similarity transformations which allows, for example, the matching of a right and left hand which would not be possible otherwise.

3.5 The Tangent Space

Once a collection of objects has been matched into optimal Procrustes position, the mean shape can be calculated and the structure of shape variability can be explored. However, rather than work in the curved shape space, analysis is often carried out in the linearised version of the shape

space: the tangent space. The advantages of using the tangent space are that linear methods are well understood and behave well. A point of tangency (the pole) is chosen, corresponding to some reference shape such as the mean shape. The landmarks on each shape in a set are then projected onto the plane which touches the shape space at this pole. These are referred to as the tangent space coordinates. If most of the objects in a dataset are close in shape then it can be shown that using the Euclidean distance in the tangent space will be a good approximation to the shape distance in the shape space (Dryden and Mardia (1998)). Hence a practical approach to exploring shape variability is to project the shapes into the tangent space and perform multivariate analysis, such as Principal Components Analysis, in a linear space.

3.6 Principal Component Analysis

Principal Component Analysis (PCA) is a multivariate statistical technique which finds a small number of uncorrelated linear combinations of the variables in a data set which explain most of the variation in the data. The general procedure is to compute the correlation or covariance matrix and use this to perform an eigen-analysis. This provides a new set of variables, or components, made up of uncorrelated linear combinations of the original variables. The resulting eigenvalues give the variance explained by each component and the eigenvectors give the loadings of each component which can be examined to establish appropriate interpretation for the components. By construction the components are independent and the total variance is preserved under the principal component transformation. The components can then be investigated in order to obtain a meaningful interpretation and explore the areas of greatest shape variation or whether there are distinct groups in the data. In conclusion, an appropriate statistical method for analysing shapes is to use Procrustes matching to align the shapes and min-

imise differences between them. The shapes are projected onto the tangent space using the mean shape as the pole, and PCA is carried out in this linear space.

3.7 Summary

Statistical shape analysis requires specialised theory in order to deal with the high dimensionality of the data. This chapter outlined methods for measuring shape and the distance between shapes. Once we have defined shape and distance, Procrustes methods can be used to standardise 3-dimensional configurations and PCA provides a new set of variables which can be examined to establish the areas of greatest shape variation.

Chapter 4

Curve Identification

4.1 Curvature

Curvature can be defined as the way in which the normal to a surface changes from point to point. There are a number of ways to measure curvature, such as the widely used Gaussian curvature K and mean curvature H. Both of these functions make use of the principal curvatures k_1 and k_2, which are the maximum and minimum directional curvatures of a surface at a given point. Together they fully describe the type of shape being measured. Goldfeather and Interrante (2004) give the following summary on how to calculate the principal curvatures. Let p be a point on a surface S. Then N_p is the unit normal to the surface S at point p and $X(u, v)$ can be considered a local parameterisation of S in a neighbourhood of p. Now we have all the parameters that are needed to completely specify our surface S and define a local coordinate system using $X_u(p)$, $X_v(p)$ and N_p, details of which are to follow in Section 4.2. We define the Weingarten curvature matrix as follows:

$$W = \begin{pmatrix} \frac{eG-fF}{EG-F^2} & \frac{fE-eF}{EG=F^2} \\ \frac{fG-gF}{EG-F^2} & \frac{gE-fF}{EG-F^2} \end{pmatrix}$$

where

$$e = N_p \cdot X_{uu}(p) \qquad E = X_u(p) \cdot X_u(p)$$

$$f = N_p \cdot X_{uv}(p) \qquad F = X_u(p) \cdot X_v(p)$$

$$g = N_p \cdot X_{vv}(p) \qquad G = X_v(p) \cdot X_v(p)$$

Carrying out an eigenanalysis on the Weingarten matrix W results in the eigenvectors k_1 and k_2, the principal curvatures required. Then the Gaussian curvature K and mean curvature H are calculated by:

$$K = k_1 \times k_2 \qquad H = \frac{k_1 + k_2}{2}$$

In other words, the Gaussian curvature K is the product of the principal curvatures and the mean curvature H is the arithmetic average of the principal curvatures.

4.2 The Shape Index

Koenderink and Doorn (1992) argue that neither the Gaussian or the mean curvature describe the "local shape" adequately. The shape index proposed in their paper describes the curvature of a shape independently of size and is expressed in a single value between -1 and 1. This shape index corresponds to a simple colour scale which can be easily interpreted by a viewer. Koenderink and Doorn (1992) also suggest a measure called curvedness which is a positive number representing the degree of curvature in the shape.

In order to represent the surface of an object, position and orientation must be removed. This can be achieved by placing the origin of the coordinate system at a chosen reference point on the shape and taking the z-axis to point along the outward normal of the surface. The following finite Taylor series then describes the deviations of the x and y coordinates:

$$z = \frac{1}{2}(k_1 x^2 + k_2 y^2), \qquad k_1 \geq k_2$$

We refer to k_1 and k_2 here as the principal curvatures which together fully describe the type of shape being measured. If $k_1 = k_2$ then we have an umbilical point, and more specifically if $k_1 = k_2 = 0$ we have a flat point. $k_1, k_2 > 0$ indicates a concave shape, whereas $k_1, k_2 < 0$ indicates a convex shape and if $k_{1,2}$ have different signs then the shape is a saddle. These principal curvatures are then combined together in a ratio to calculate the shape index using the following formula:

$$s = \frac{2}{\pi} tan^{-1} \frac{k_2 + k_1}{k_2 - k_1}$$

The resulting value of s will lie between -1 and 1, with convex and concave shapes at either end of the scale. The most extreme convex shape, a cap, is represented by a value of 1 and the most extreme concave shape, a cup, is represented by a value of -1. Other complementary pairs that fit together like a stamp into its mould also have opposite signs. The cylindrical shapes, a ridge and a rut, are found at 0.5 and -0.5 respectively and a symmetrical saddle is indicated by a value of 0. In order to obtain a quick and easy visual identification of a shape type, the shape index can be mapped to a colour scale. This allows a surface to be coloured in according to the local shape. The colour scale used by Koenderink and Doorn (1992) was determined according to the following properties:

- Convexities, concavities and saddle shapes should be immediately distinguishable.
- Shapes that are easy to identify should be represented by recognizable colours.
- The extreme ends of the scale (convex vs concave shapes) should stand out against each other.

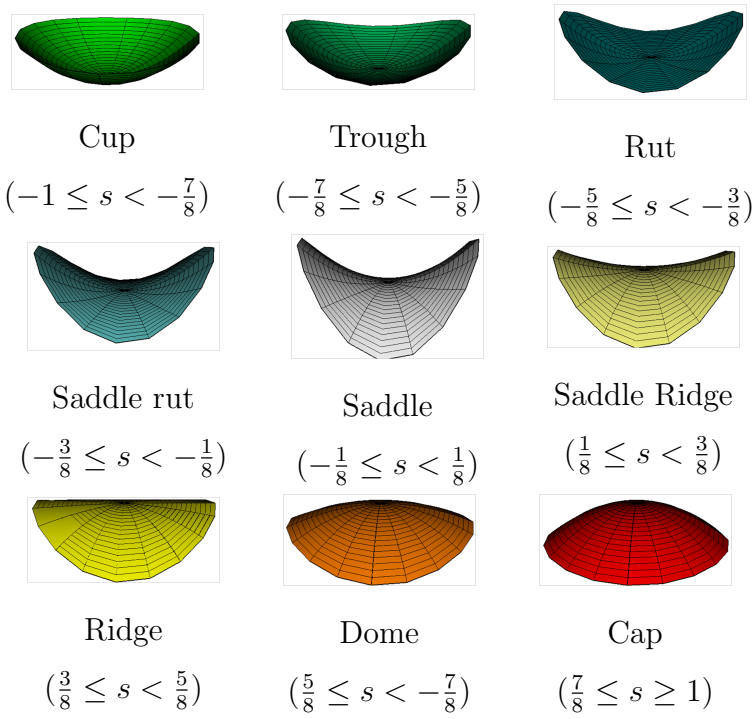

Figure 4.1: Shape types with corresponding shape index values and colour coding

- Complementary shapes (e.g. cap and cup) should be assigned complementary colours.

These properties lead to the scale shown in Figure 4.1.

Once the nature of the local shape has been determined and coloured accordingly, a measure of the strength of the surface curvature can be calculated. This is called curvedness and is calculated using the following formula:

$$c = \sqrt{\frac{k_1^2 + k_2^2}{2}}$$

The curvedness is geometrically meaningful since it is independent of the surface's position in space. Koenderink and Doorn (1992) suggest using a monochrome scale to visually represent curvedness, with darker shades indicating more intense shape curvature.

The paper by Cantzler and Fisher (2001) compares the mean and Gaussian curvature functions to the shape index proposed in Koenderink and Doorn (1992). They conclude that, although the two methods are extremely

similar, Koenderink's approach is preferable overall. It is more stable at low thresholds and is advantageous when dealing with complex shapes and image noise.

4.3 B-splines and Penalised Splines

A spline is a smooth polynomial that is defined by multiple subfunctions and has a desired degree of smoothness at the points where the subfunctions join together. A basis spline, or B-spline, is a spline function that is constructed by joining polynomials together at fixed points. These connection points are referred to as knots. In other words, the range x_{min} to x_{max} is divided into m' equal intervals by $m'+1$ knots and each interval is covered by $q+1$ B-splines of degree q. Hence the total number of knots required to form the B-splines is $m'+2q+1$ and the number of B-splines in the regression is $m = m' + q$. Let $B_j(x;q)$ denote the value of the jth B-spline of degree q at x for a given set of equally spaced knots. Then a fitted curve $\hat{y}$ to data (x_i, y_i) is the linear combination:

$$\hat{y}(x) = \sum_{j=1}^{m} \hat{\alpha}_j B_j(x;q)$$

The choice of the number of knots is critical when modelling with B-splines as too many knots results in overfitting of the data whereas too few knots results in underfitting of the data. In order to avoid this problem of over-fitting an additional penalty function can be incorporated in the calculation of spline functions. Penalised splines, or P-splines, are a combination of B-splines and difference penalties on the estimated coefficients of the fitted curve (Eilers and Marx (1996)). They have no boundary effects, are a straightforward extension of linear regression models, conserve moments of the data and have polynomial curve fits as limits. Consider the regression of n data points (x_i, y_i) on a set of m B-splines $B_j(.)$. The least squares objective function to

minimise is:

$$S = \sum_{i=1}^{n} \{y_i - \sum_{j=1}^{m} \alpha_j B_j(x_i)\}^2$$

Let the number of knots be sufficiently large that the fitted curve shows more variation than is justified by the data. In order to make the resulting curve less flexible, Eilers and Marx (1996) propose to base a penalty on finite differences of the coefficients of adjacent B-splines:

$$S = \sum_{i=1}^{n} \{y_i - \sum_{j=1}^{m} \alpha_j B_j(x_i)\}^2 + \lambda \sum_{j=k+1}^{m} (\Delta^k \alpha_j)^2$$

where $\Delta\alpha_j = \alpha_j - \alpha_{j-1}$ and, for example, $\Delta^2\alpha_j = \alpha_j - 2\alpha_{j-1} + \alpha_{j-2}$ and so on. This reduces the dimensionality of the problem to m, the number of B-splines, instead of n, the number of observations. The parameter λ allows for continuous control over the smoothness of the fitted curve.

4.4 Principal Curves

Principal Component Analysis (PCA) is a multivariate statistical technique which finds a small number of uncorrelated linear combinations of the variables in a data set which explain most of the variation in the data. However, if the data we are dealing with is non-linear then this method is not appropriate. Principal curves (Hastie and Stuetzle (1989)) allow us to summarise non-linear data by fitting a smooth curve that runs through the centre of a data set. In order to calculate a principal curve we introduce the idea of self-consistency. Given a matrix X of n observations on p variables in $\mathbb{R}^p$, we define $f(\lambda)$ to be a one-dimensional curve, where λ indicates the positioning of points along the curve parameterised by arc length. The projection index, $\lambda_f(x)$, provides the value of λ for which the curve $f(\lambda)$ is closest to a particular point x. A curve f is self-consistent if

$$\mathbb{E}(X|\lambda_f(X) = \lambda) = f(\lambda)$$

i.e. for every point on the principal curve, if we take the data that project onto it and average them, then this should match the chosen point. In practice, this step of the principal curve calculation is achieved using a smoothing procedure which calculates the local averages along the curve. The analysis in the project has used P-splines (see Section 4.3) to carry out the smoothing step of the principal curve algorithm.

The general algorithm for calculating a principal curve is to select an initial curve (often a straight line such as the largest principal component) and to check if this curve is self-consistent by projecting and averaging points. If it is not self-consistent then the process is iterated using the new curve obtained through the projection and averaging. If it is self-consistent then the curve found is a principal curve.

4.5 Local Principal Curves

There are some limitations with the general principal curve method, such as a dependence on the chosen starting line and problems of bias. Furthermore principal curves are unable to handle branched or disconnected curves. In order to solve these problems Einbeck *et al.* (2005) propose an algorithm for local principle curves (LPC), which examines only a local neighbourhood around each point under consideration. The LPC algorithm is as follows:

1. Choose an appropriate starting point x_0 and define $x = x_0$.

2. Calculate the local centre of mass around x, $\mu(x)$. The local centre of mass around x is essentially a weighted average of the local neighbourhood of points and is defined by the following formula:

$$\mu(x) = \frac{\sum_{i=1}^{n} K_H(X_i - x)X_i}{\sum_{i=1}^{n} K_H(X_i - x)}$$

   where K_H is a d-dimensional kernel function and H is a bandwidth matrix.

3. Perform a PCA locally at x and find the first local principal component, denoted γ.

4. Follow γ from x for some previously chosen distance t_0. This provides the new value of x i.e. $x = \mu(x) + t_0\gamma$.

5. Iterate the process until $\mu(x)$ remains constant.

The starting point chosen in step 1 could be chosen at random or based on a density estimate. Einbeck *et al.* (2005) suggest that the first approach should be preferred since it is computationally easier and handles crossings better. The process will naturally remain constant when the end of the data points is reached, however the algorithm can be stopped earlier than step 5, for example when the difference between consecutive curves is less than a chosen threshold value.

Some additional steps can be included in the above algorithm in order to ensure that all types of curves can be handled. One end of a curve may have been ignored due to the fact that we follow the first principal component along only one direction. Taking the opposite direction would have been equally as valid and hence the following step is incorporated:

6. For the opposite direction of γ, $-\gamma$, perform steps 4 and 5.

Finally, if the curve is made up of several disconnected branches then it will be necessary to have multiple starting points:

7. For a set of starting points S_0 choose, without replacement, a new starting point $x_0 \in S_0$ and start the process again from step 1.

This step will ensure that the algorithm does not fall short of fully describing the data. The choice of starting points can be chosen by hand in order to ensure each section of the curve is covered, or can be chosen at random, which may be computationally easier.

4.6 The Lips

The development of an automatic curve identification algorithm was focused on the lips. Given a set of points representing the lips on a face, an estimate of the three key curves (the upper and lower lips, and the midline valley where the upper and lower lips meet) can be obtained by fitting a curve through the candidate points. As discussed in Bowman *et al.* (2012) the upper and lower lip boundaries can be problematic as fitted curves often wrongly track other ridges in the skin shape around the sides of the mouth. However, the midline where the lips meet is well defined and provides a solution through the identification of the points at the corners of the mouth. Any points that lie beyond these corner points can be discarded and the upper and lower lip curves can be constrained to meet, and end, at these locations.

Accurately finding the corner points of the mouth is essential in order to give an exact representation of the midline curve and to precisely tie down the upper and lower lip curves. Finding the corner points is a complicated task and a number of methods were explored. For instance, the corner points are characterised by very strong curvature where the midline curve rises to join the surrounding skin. A principal curve can be fitted to the midline, parameterised by $\{x(s), y(s), z(s)\}$ where s is the arc length along the fitted curve. The curvature of this curve can then be calculated using the following formula, based on the one given by Koenderink and Doorn (1992) :

$$\frac{\sqrt{\{x''y' - y''x'\}^2 + \{x''z' - z''x'\}^2 + \{y''z' - z''y'\}^2}}{(x'^2 + y'^2 + z'^2)^{2/3}} \tag{4.1}$$

The resulting curvature is a function of the arc length s and the corner points of the mouth can now be identified as the first and last local maxima of this function, denoted by (x_l, y_l, z_l) and (x_r, y_r, z_r) respectively. Alternatively, the shape index can be used in order to identify the corner points. In this case, areas with the most extreme concave shape, a cup, can be used as opposed to points with the strongest curvature value. If the two areas consisting of

Figure 4.2: Identifying the cheilions using the automatic curvature method

coordinates with a shape index value between $-\frac{7}{8}$ and -1 (a cup) can be identified along the midline section of the lips, and the centre of these areas found, then the resulting points can be considered to be the corner points of the mouth.

Both of these methods were explored, but it was considered that neither gave sufficiently accurate results, in comparison to manual anatomical landmarking. Figure 4.2 shows an example of cheilions found using the first method described above (the first and last local maxima of the curvature of a principal curve along the midline). The left-hand cheilion seems to be very close to where we would manually place the landmark, although it is not perfectly accurate. However, the right-hand cheilion is clearly not identified appropriately using this method. The function has found a local maxima that is far from the lips themselves, and has classified this as the second cheilion.

Since the anatomical landmarks had already been marked up for the 12 years olds, and were also available for all the previous time-points, it was decided that it would be sensible to make use of this information. Landmark analysis is a widely used method and provides suitably precise results. Figure 4.4 shows the isolated midline section of the lips with the anatomical landmarks at the corners of the mouth, also referred to as the cheilions (see 2.2). Once the area of interest and the anatomical landmarks have been identified a curve can be fitted through the three-dimensional coordinates,

with the added restriction that the curve must lie between the two chosen landmarks. The choice of curve to be fitted through the points must now be made, and the following sections describe two methods which can be used. Section 4.6.1 outlines how the shape index and principal curves can be used to create lip curves. Section 4.6.2 introduces the concept of planepath curves and describes how they can be used to create a curve between two points of interest.

4.6.1 Using the Shape Index

The shape index proposed by Koenderink and Doorn (1992) can be used to identify key features of the face. Figure 4.3 exemplifies this using an image captured for one of the 12 year olds. Figure 4.3 shows the subject's face colour coded by the shape index. We can see that the valley of the midline of the lips, the ridges of the nose, the brow and the upper and lower lips, and the deep cups at the inner eyes have all been highlighted by the shape index. Hence, the information obtained from the shape index can be used in the automatic identification of curves as we can isolate areas of the face according to their shape type and then fit a principal curve through these particular points. For example, if we threshold the shape index between $-3/8$ and -1 then we obtain the image shown in Figure 4.3. The shape type has been restricted to a rut, trough or spherical cup and the remaining features of the face include the midline and the eye cups.

In the case of the lips, we begin by isolating the midline of the lips, identifying the cheilions and fitting a principal curve between these two points, as shown in Figure 4.4. Further restrictions can be incorporated into the principal curve algorithm, so that the curve passes through additional available landmarks such as the middle point of the lips, the stomion. Such constraints can conveniently be incorporated into the principal curve algorithm at the smoothing step, using p-splines. As detailed in Bowman *et al.* (2013), a p-spline curve can be expressed in the same form as a linear regression,

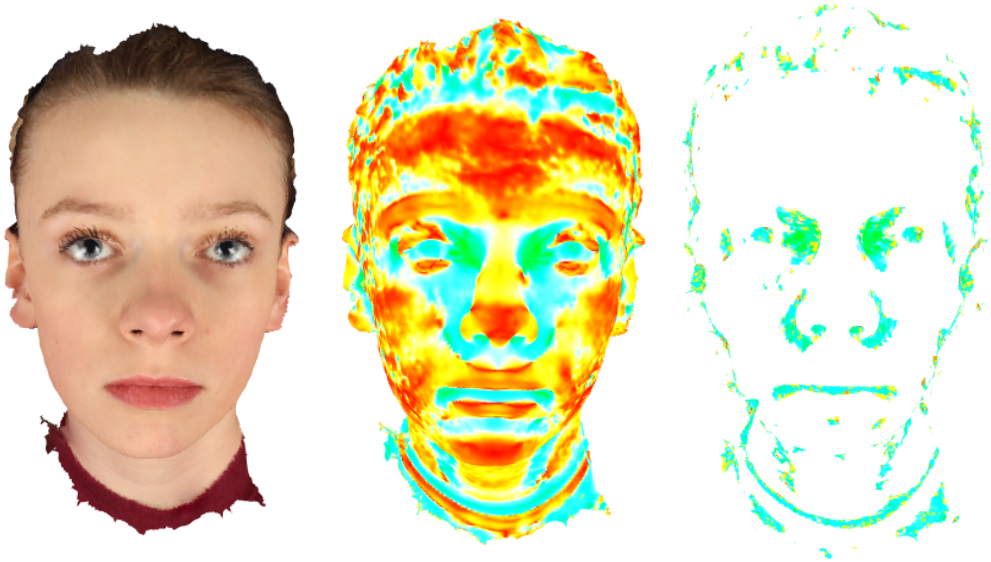

Figure 4.3: Using the shape index to identify key features of the face of a 12 year old child, with the natural image (left), the shape index (middle) and the shape index thresholded between $-\frac{3}{7}8$ and -1

$x = B\beta$, where B is a design matrix which evaluates a set of local B-spline basis functions at the values of the observed covariate. The model can be fitted to the observed set of data by using the regression coefficients $\hat{\beta}$ which minimise the penalised sum-of-squares:

$$S(\beta) = (x - B\beta)^T(x - B\beta) + \lambda\beta^T D_2^T D_2\beta$$

where D_2 is a matrix of the second differences of the elements of β. If we require that the curve passes through a specific point then we use the constraint $A\beta = c$, where the columns of the matrix A evaluate the basis functions at the constraint locations and the vector c contains the constrained response values. Now the sum-of-squares $S(\beta)$ can be minimised with respect to the constraint $A\beta = c$ and the constrained coefficients $\hat{\beta}_c$ are defined by:

$$\hat{\beta}_c = \hat{\beta} + (B^T B + D_2^T D_2)^{-1} A^T [A(B^T B + D_2^T D_2)^{-1} A^T]^{-1}(c - A\hat{\beta}).$$

Additionally, weights can be incorporated into the principal curve algorithm. These weights are based on the appropriate value of the shape index

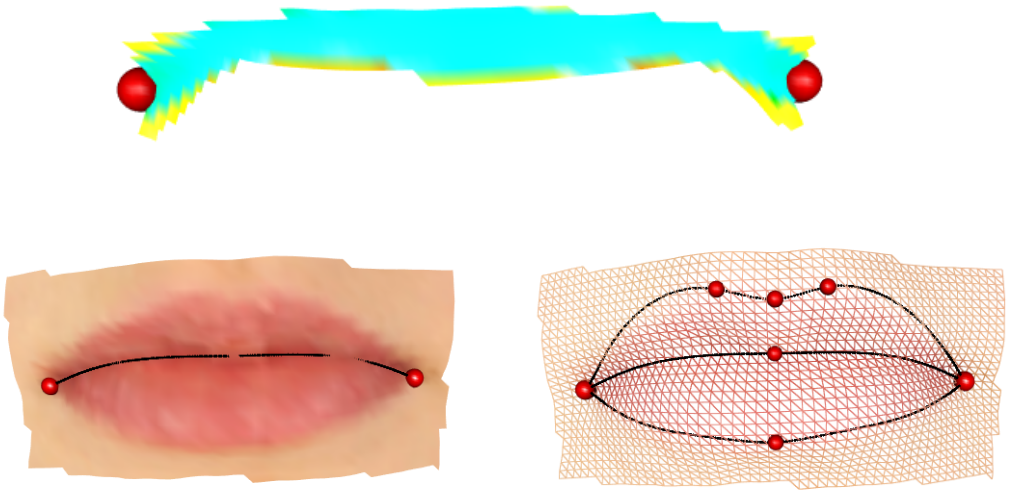

Figure 4.4: Isolated midline of the lips with cheilions (top), midline curve (bottom left) and midline, upper and lower curves, with anatomical landmarks (bottom right)

for the shape type of interest - for example, in the case of the midline the weights are centred on $-3/8$, a rut shape, and take the form:

$$weights = \exp(-0.5 \times (\theta_m + 3/8)^2 \times 8^2)$$

where θ_m is a vector of the shape index values along the midline curve.

The same method is used for the upper and lower lips, using the appropriate anatomical landmarks for additional restrictions on the principal curve and the appropriate values of weights. Specifically, the upper lip curve is forced to pass through the left and right cheilions, the left and right crista philtri and the labiale superius, and the lower lip curve is forced to pass through the left and right cheilions and the labiale inferius. Both the upper and lower lips are characterised by a distinct convex shape and hence the weights applied in these cases are centred on $5/8$, a dome shape, and take the form:

$$weights = \exp(-0.5 \times (\theta_l - 5/8)^2 \times 8^2)$$

where θ_l is a vector of the shape index values along either the upper or lower lip curve. The final result is shown in Figure 4.4. We can see that the curves follow the natural curves of the mouth reasonably well, guided by the appropriate anatomical landmarks.

There are still refinements that could be made to this algorithm. For example, one key issue is that the curves do not necessarily lie on the facial surface and therefore are not always as precise as would be desired. However, methods for overcoming this problem are not straightforward and so, for the purposes of this project, the current algorithm was considered satisfactory for analysis. The algorithm should also be able to find curves on other areas of the face, such as the nose profile and the brow ridge, but again in the context of this project only the lip curves have been explored.

4.6.2 Using Planepath Curves

An alternative approach to automatically identifying curves is to calculate the shortest distance along the facial surface between two points of interest, referred to as a planepath curve. The planepath curve algorithm is as follows:

1. Identify two points of interest: x_1 and x_2.
2. Find the normals at these points, denoted n_1 and n_2.
3. Take the average of these normals, denoted n.
4. Find the plane which contains x_1 and x_2.
5. Find the set of points between x_1 and x_2 which lie on the edges of the mesh and which lie on the plane.
6. Repeat steps 4 and 5 for a range of angles which tilt the plane about the normal direction, n.
7. Choose the path which has the smallest length to be the final planepath curve between x_1 and x_2.

An example of the planepath curve is shown in Figure 4.5. The planepath curve describes the shape of the upper lip reasonably well, although the finer detail of the shape is not captured. The main problem is that the planepath is too rigid and does not follow the arch of the upper lip, particularly between the left cheilion and left crista philtri, and the right cheilion and right crista philtri. However, it is clear that overall the planepath does give a good representation of the upper lip shape and importantly it sits on the surface of the shape. This is not the case when using principal curves, as discussed in Section 4.6.1. Furthermore, the planepath curve can be applied to any region of the face, provided there are accurate landmarks available. Figure 4.5 demonstrates how the planepath curve can be used to provide a comprehensive description of the central features of a face with sufficient accuracy.

4.7 Summary

There are a number of ways to describe the curvature of a surface, including Gaussian and mean curvature. The shape index proposed by Koenderink and Doorn (1992) provides an alternative measure of shape which is independent of size and can be easily visually recognised using a colour scale. A method of curve identification was developed using a combination of the shape index and principal curves, a method which summarises non-linear data. This method was compared to an existing approach which simply calculates the shortest distance along the facial surface between two points of interest, the planepath curve. The planepath curve gives a reasonably accurate representation of the facial shape and can be applied to any area of the face. Therefore, despite being slightly restrictive, planepath curves are preferable to principal curves in their current form. It is likely that further development of the principal curve method may produce curves of extremely high quality.

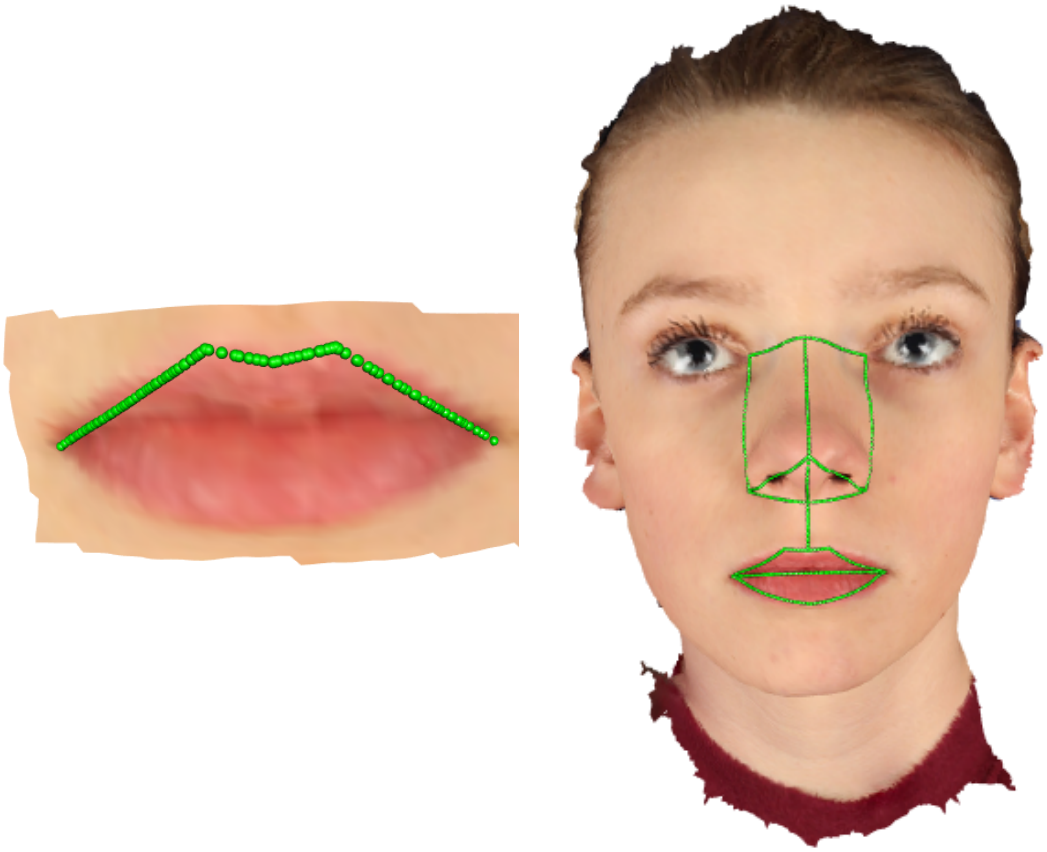

Figure 4.5: Using planepath curves to identify key features of the face of a 12 year old child, with the upper lip (left) and the full face (right)

Chapter 5

Longitudinal Analysis

5.1 Data for Longitudinal Analysis

As discussed in Section 2.4, the final number of 12-year images which were suitable for analysis was 35 (of the 39 images captured, 2 children had their lips parted due to braces and 2 images were not of high enough quality around the central features of the face). This gave a total of 203 images to include in longitudinal analysis - including the previous images taken of the 35 final subjects. A summary of the images used in the longitudinal analysis is shown in Table 5.1.

Table 5.1: Summary of the Images Used in Longitudinal Analysis

Time point (months)	3	6	12	24	60	144
Number of images	29	34	35	35	35	35
Mouth Closed	6	6	6	5	35	35
Mouth Open	23	28	29	30	0	0

Table 5.1 shows that the most missing data occurs at the first time point when the infants were 3 months old. The most likely reason for this is that these subjects may not have been recruited until after the 3 month mark, making the 6 month time point their initial visit. However, the proportion of

missing data is not high - there are 203 available images of a potential 210.

Some key factors which can effect the data collected from infants, which was discussed in Section 1.3, are the child's age, the level of cooperation given by each child and the length of the examination session. Although the C3D camera system allows images to be captured in a matter of seconds, some of the images were taken while the subjects were still very young and therefore would clearly be unable to follow specific instructions regarding their placement and facial expressions. In most image sessions the subject will be asked to keep their lips closed, in order to capture comparable pictures which reflect the natural shape of the mouth. For many infants this specific facial expression was not possible. In some cases multiple images were captured in order to achieve the correct facial expression. For the purposes of analysis, those subjects who had multiple images available have had the image with their lips closed selected, rather than the image with their lips open. However, for other subjects multiple images were not taken, and therefore a number of subjects between the ages of 3 and 24 months have images with open mouths, as detailed in Table 5.1. In order to accurately represent the face shape and allow for comparisons between subjects and between time points, it was decided that only the region above (and including) the upper lip would be considered in the longitudinal analysis.

The method used to represent the face curves of the 203 subjects is the planepath curve, which is discussed in detail in Section 4.6.2. Although using principal curves (Section 4.6.1) may provide a more accurate and flexible curve, the method still requires some fine-tuning and was not ready to be applied to other regions of the face, apart from the lips. Planepath curves are able to provide an accurate representation of the central features of the face and a final set of 13 curves were chosen to provide this shape summary (see Figure 5.1). Since the distances between the labiale superius and the right and left crista philtri are so small, particularly in the early images, the three landmarks are considered to provide adequate information, instead of an

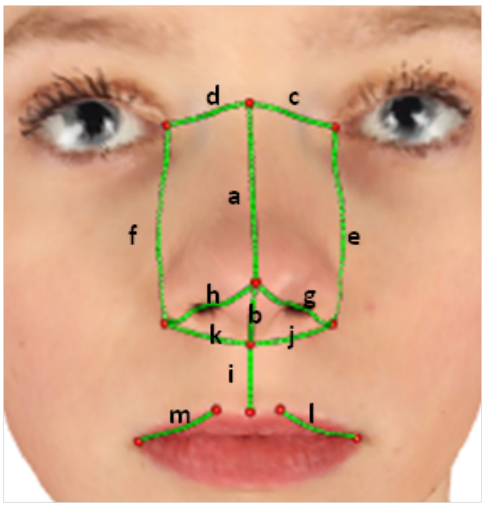

Figure 5.1: Planepath curves and landmarks used in longitudinal analysis (a: midline nasal profile; b: columella; c: nasal root left; d: nasal root right; e: nose lateral left; f: nose lateral right; g: nose ridge left; h: nose ridge right; i: philtrum; j: nasal base left; k: nasal base right; l: upper lip left; m: upper lip right)

additional two planepath curves. Each curve was assigned a specific number of points depending on its size, which was to be consistent across all cases. The 12 landmarks used in addition to the 13 curves are the sellion, pronasale, subnasale, right and left endocanthions, right and left alare crests, labiale superius, right and left cheilions and right and left crista philtri (see Section 2.4). The final array of data, combining the 13 curves and 12 landmarks, consists of 87 three-dimensional points for each of the 203 images.

5.2 Procrustes Matching and Principal Components Analysis

Before longitudinal analysis can be carried out, the shapes must first be standardised so that they are matched in rotation and centred at a common origin using the method of Procrustes matching (see Section 5.2 for an explanation of this method). Since we are interested in growth, it would not

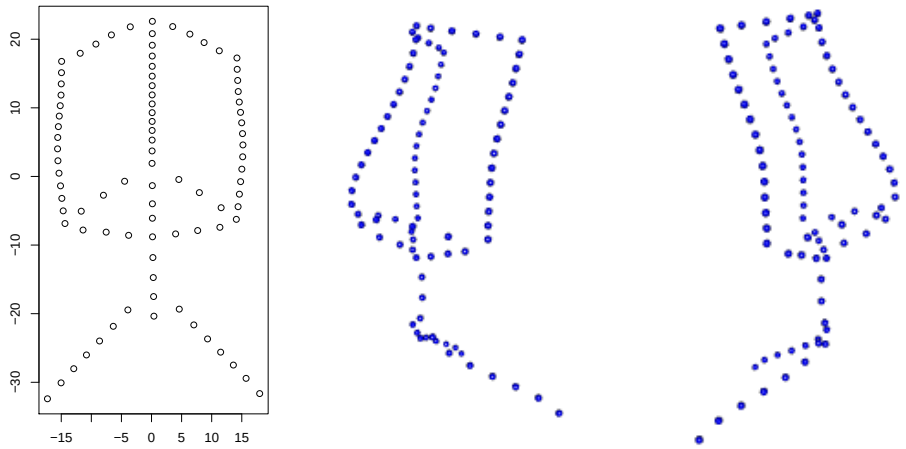

Figure 5.2: Mean shape viewed from the front (left), left-hand side (middle) and right-hand side (right)

be appropriate to include the scaling of the images so this step was removed from the usual Procrustes matching process. The mean shape resulting from the Procrustes matching process can be seen in Figure 5.2.

Now that the shapes have been optimally matched according to rotation and placement, a principal components analysis is carried out. Figure 5.3 shows the first and second principal components from this analysis, with sex and then time point identified. These plots suggest that there may be no clear difference in facial shape between the sexes according to either of the first two components. However, they indicate that there may be a difference in facial shape depending on the time point, a separation that is distinguished by the first principal component only. Although there is some overlap between the 3, 6, 12 and 24 month time points, a gradual decrease in the first principal component can be identified as the subjects get older, and the separation between the 60 and 144 month time points is considerably more evident.

The principal components analysis returns 203 principal components, which must be reduced to a specified number of components for further analysis. It was considered that enough principal components were required

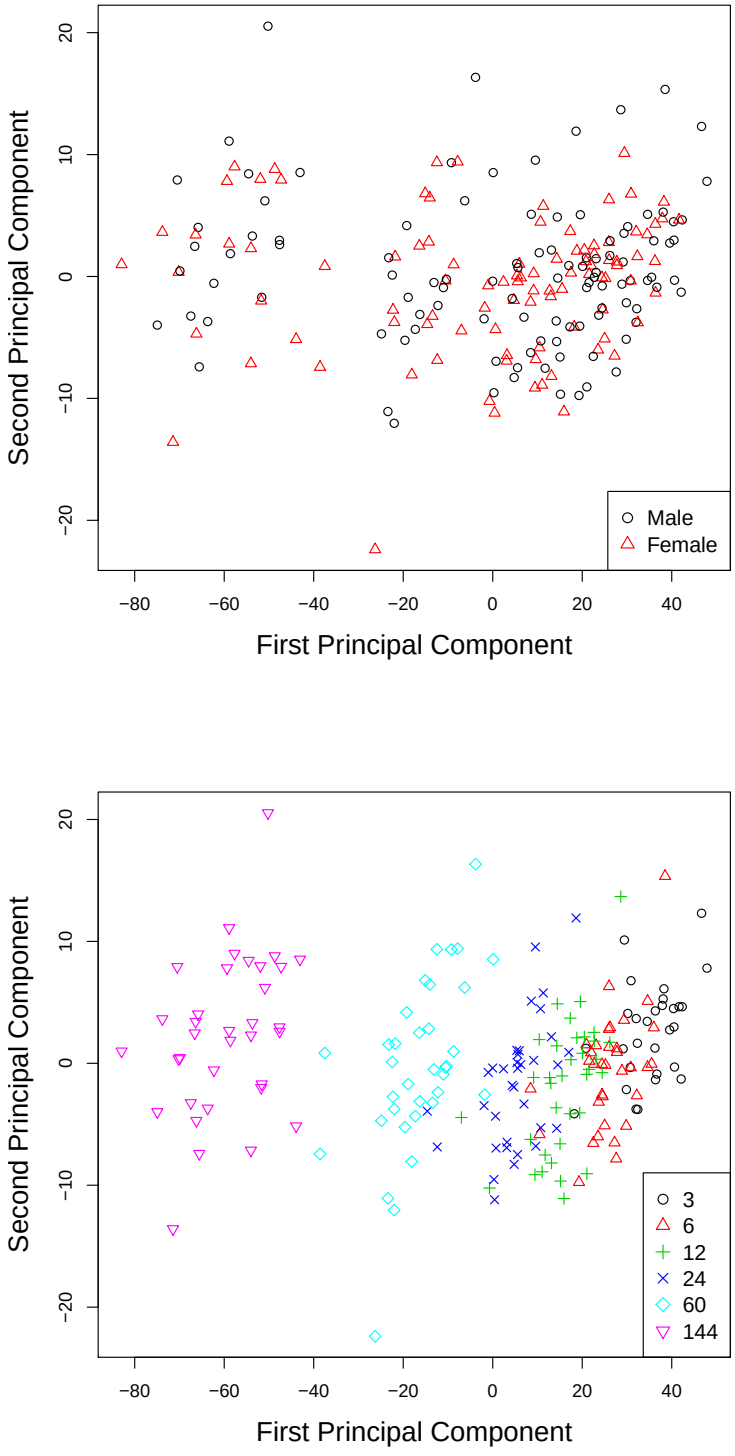

Figure 5.3: Plots showing the first principal component against the second principal component split by sex (top) and time point (bottom)

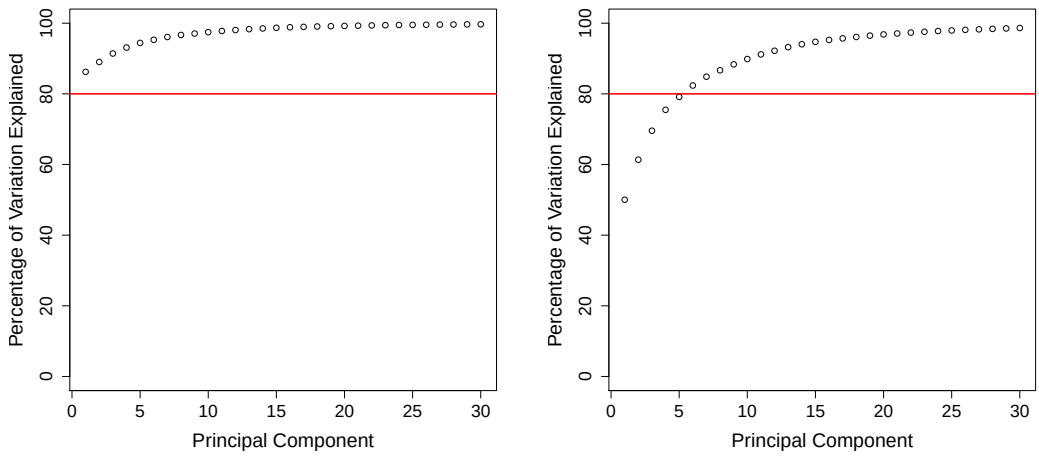

Figure 5.4: Scree plots showing the first 30 principal components: without scaling (left) and with scaling (right) included in the Procrustes matching process

in order for 80% of the variation in the data to be explained. However, the first component explains 86.23% of the variation, an extremely large amount which is already greater than the required level of 80% (see Figure 5.4). Therefore, it was not clear how many further components should be retained. The Procrustes matching and principal components analysis was repeated with scaling now included. This allowed the components to be inspected without the first component dominating. Figure 5.4 shows the scree plot for the second analysis process. We can now see that the first 6 components capture over 80% of the variation in the data. The exact value of cumulative variation explained by the first 6 components is 82.36% (see Table 5.3). Therefore, in the next stages of the analysis 6 principal components were used to represent the data.

The 6 chosen scores were plotted against time in order to explore any patterns in the data. Figure 5.5 shows a decreasing trend in the first principal component over time - as time increases, the value of the first principal curve decreases. This pattern appears to be the strongest out of all the 6 principal components. The other boxplots in Figure 5.5 do not suggest that there are

Table 5.2: Percentage of Variation Explained by the First 6 Principal Components (scaling not included in PCA)

Principal Component	1	2	3	4	5	6
Variation Explained (%)	86.23	2.79	2.39	1.68	1.32	0.88
Cumulative Variation (%)	86.23	89.03	91.42	93.10	94.42	95.31

Table 5.3: Percentage of Variation Explained by the First 6 Principal Components (scaling included in PCA)

Principal Component	1	2	3	4	5	6
Variation Explained (%)	50.03	11.34	8.20	5.92	3.65	3.23
Cumulative Variation (%)	50.03	61.37	69.57	75.48	79.14	82.36

any other strong relationships between the next five principal components and time.

It was proposed that the first principal component may represent size. Figure 5.6 plots the inverse of the first principal component against the centroid size (defined in Section 3.3) and shows an extremely strong linear correlation between the two variables. This illustrates that the first principal component does represent the shape size - as the principal component increases the shape size decreases.

We can also determine the shape change represented by each principal component by comparing the extremes of each component, as displayed in Figure 5.7. It should be noted that the first principal component explains a very large proportion of the variation in the data, but the next five components explain very little of the variation (see Table 5.2). This means that the latter five components may not have particularly strong or meaningful interpretations. Figures 5.7(a) and (b) confirm that the first principal component explains the variation in size change, with a very large and clearly distinguishable pattern. Figures 5.7(c), (d), (e) and (f) suggest that the second principal component primarily contains information about the elongation

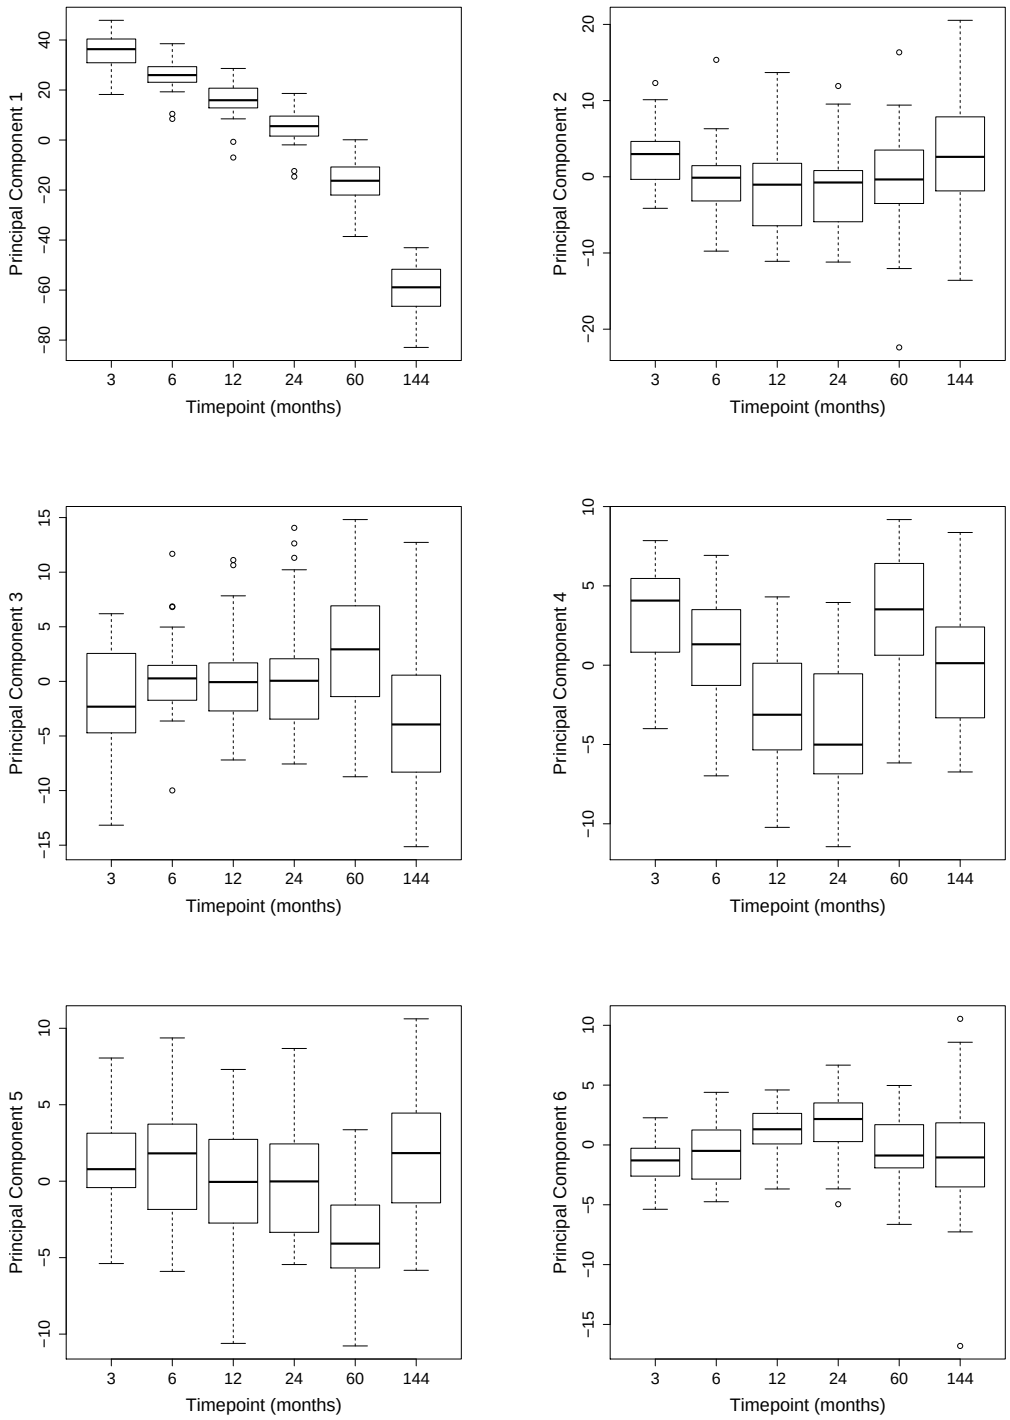

Figure 5.5: Boxplots showing each of the first 6 principal components against time

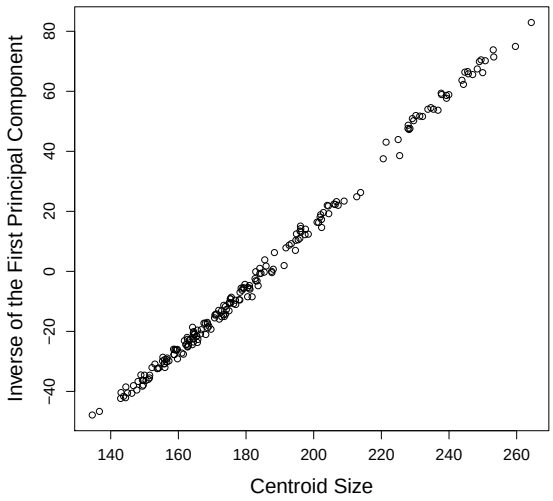

Figure 5.6: Plot showing the centroid size against PC1

of the philtrum and mouth area in both upward and downward directions, whereas the third principal component represents a broadening of the nose area. The fourth and fifth principal component contain similar shape information, which mostly captures a change in the angle and size of the upper lip (see Figures 5.7(g), (h), (i) and (j)). In particular, the fourth principal component shows a small change in the amount of protrusion of the upper lip. A slight difference between the information captured by these components is that the fourth principal component also represents an increase in the size of the tip of the nose, whereas the fifth principal component represents a broadening and flattening of the upper portion of the nose area. Figures 5.7(k) and (l) show a very weak shape change in the sixth principal component. This is unsurprising since the sixth principal component explains less than 1% of the variation in the data (see Table 5.2).

5.3 Linear Mixed Effect Models

Longitudinal data consists of a small number of measurements that are made on a large number of individuals. Here we have 6 time points at

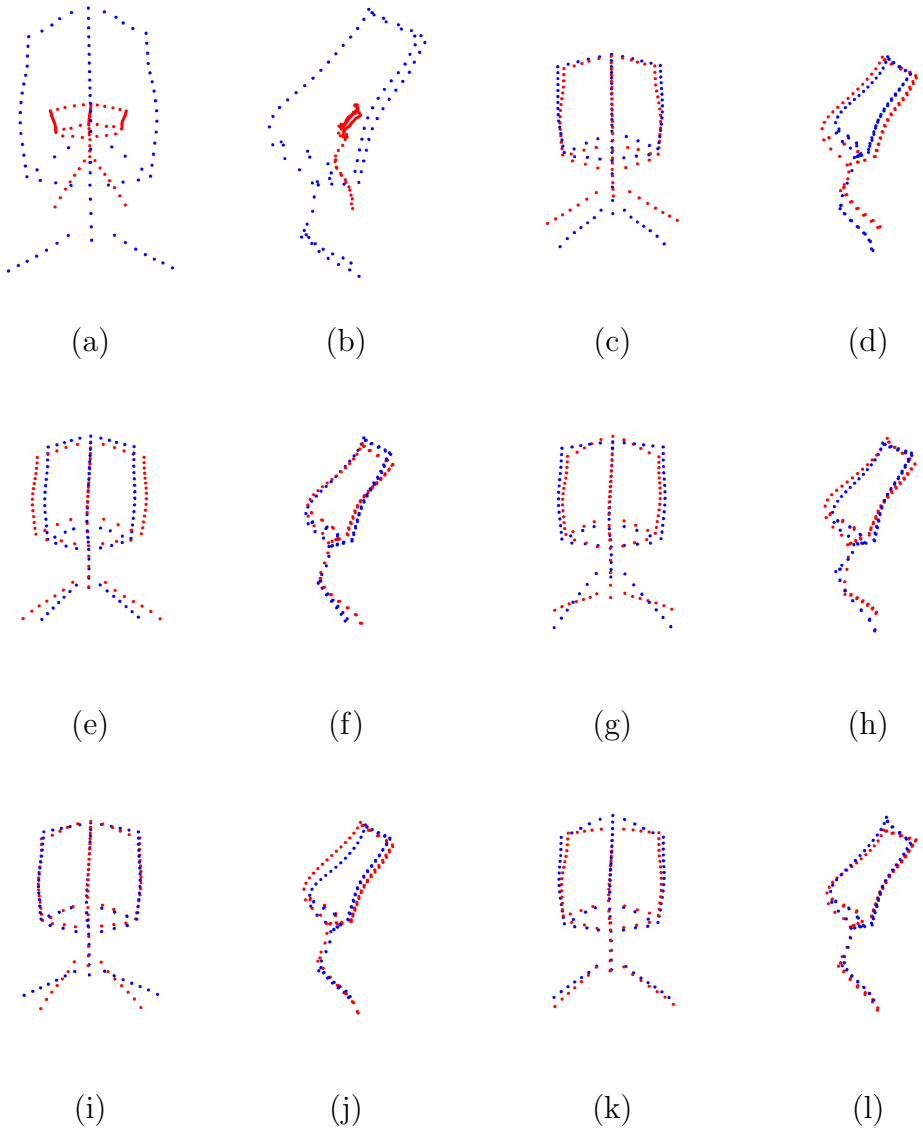

Figure 5.7: Frontal and side views of the extreme values of the principal components analysis without scaling: PC1 ((a) and (b)), PC2 ((c) and (d)), PC3 ((e) and (f)), PC4 ((g) and (h)), PC5 ((i) and (j)) and PC6 ((k) and (l)). Blue: maximum. Red: minimum.

which 35 individuals were measured. A common approach to the analysis longitudinal data is to use linear mixed effects (LME) models. The key difference between linear models and LME models is that in the LME model the error term includes a random effect component in addition to the usual measurement error term. This random effect term is required in the model since there are multiple measurements taken on each subject and hence the measurements are not independent (a crucial assumption for a linear model). In an LME model each subject is allowed to have their own deviation from an overall mean and their own deviation from the overall trend over time. Here we will model the first 6 principal components by the 6 time points, and then by sex also, allowing for a different score value for each subject at each time point. We denote the first six principal components by PC1, PC2, PC3, PC4, PC5 and PC6 respectively.

5.3.1 Basic Model

The first model (Model 1) considered was a simple linear mixed effects model of the chosen principal components modelled by the 6 time points and allowing for a different score value for each subject at each time point. The model formula is given by:

$$y_{ijk} = \mu_k + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \delta_{ik}$$

where i denotes the subject, j the time point and k the principal component, and μ_k represents the fitted mean value, α_j the additive effect of time, β_k the additive effect of each principal component and $(\alpha\beta)_{jk}$ the interactive effect of time and each principal component together. Finally, δ_{ik} is the random effects vector allowing each individual to have their own principal component score value at each time point.

We are interested in the results of the fixed effects in the model, shown in Table D.1 (for both fixed and random effects see Table D.1 in Appendix D). Figure 5.8 summarises these results with the estimated values of the fixed

effects marked on the boxplots of the principal component scores. The model estimates a strong decreasing trend in the first principal component, very close to the median value of the data. Since the first principal component represents size, this tells us that the size of the face increases over time. The next five components do not allow for very clear interpretations. This is due to the dominating nature of the first principal component and the resulting small amounts of variation explained by the other five, as discussed in Section 5.2 above. In fact, since the sixth principal component does not describe any meaningful shape change, the interpretation of its estimated values is also not meaningful. The second principal component has a weak quadratic trend in the model estimates. This suggests that the pattern of shape change contained in the second principal component (an elongation of the philtrum of the nose) reverses over time, however the reason for this is not clear. The third, fourth and fifth principal components all show either a steadily increasing or decreasing trend, with an outlier at either the 60 or 144 month time point. This seems to suggest that at a certain time point after infancy, the particular shape changes reverse. However, since these principal components explain such small amounts of variation in the data (see Table 5.2) these patterns should not be considered with too much weight.

Table 5.4: Estimated Values of the Fixed Effects in Model 1

Fixed Effect	Intercept	Time (6 months)	Time (12 months)	Time (24 months)	Time (60 months)	Time (144 months)
Estimated Value	34.894	-8.976	-18.967	-29.505	-51.306	-94.384
Fixed Effect		PC2	PC3	PC4	PC5	PC6
Estimated Value		-32.413	-36.517	-31.644	-33.508	-36.246
Fixed Effect		Time(6)*PC2	Time(12)*PC2	Time(24)*PC2	Time(60)*PC2	Time(144)*PC2
Estimated Value		5.724	14.555	25.334	48.546	94.252
Fixed Effect		Time(6)*PC3	Time(12)*PC3	Time(24)*PC3	Time(60)*PC3	Time(144)*PC3
Estimated Value		10.816	20.838	31.711	55.879	93.336
Fixed Effect		Time(6)*PC4	Time(12)*PC4	Time(24)*PC4	Time(60)*PC4	Time(144)*PC4
Estimated Value		6.414	13.258	22.436	51.395	90.873
Fixed Effect		Time(6)*PC5	Time(12)*PC5	Time(24)*PC5	Time(60)*PC5	Time(144)*PC5
Estimated Value		8.631	17.625	28.026	46.133	94.629
Fixed Effect		Time(6)*PC6	Time(12)*PC6	Time(24)*PC6	Time(60)*PC6	Time(144)*PC6
Estimated Value		9.640	21.433	32.543	52.502	95.013

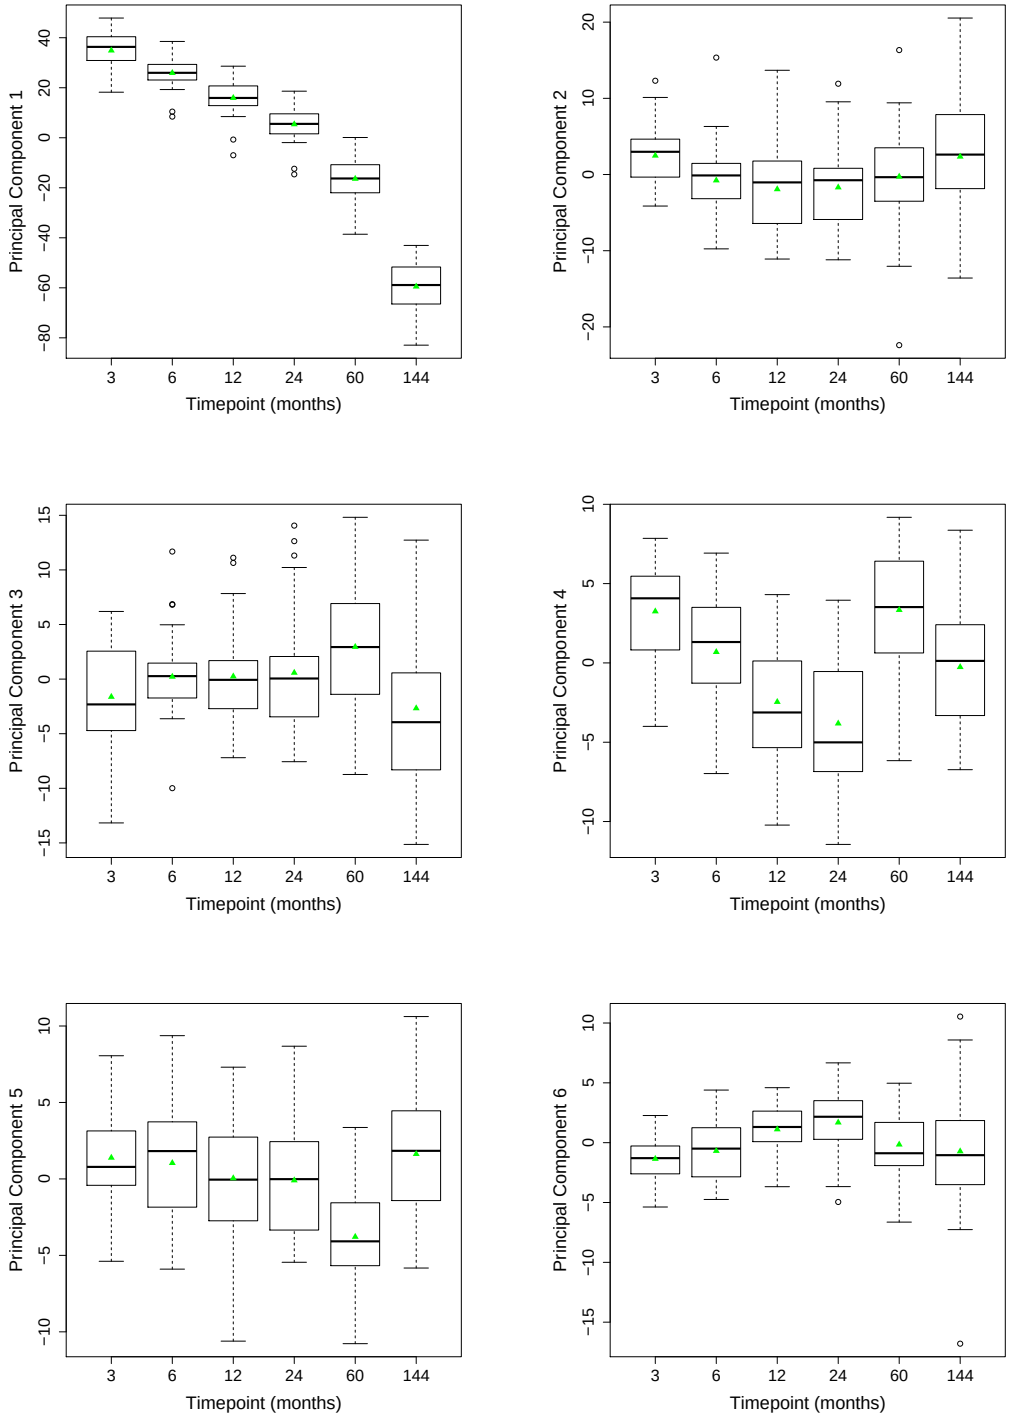

Figure 5.8: Boxplots of the first six principal components (without scaling removed) against time with estimates from Model 1 marked: PC1 (upper left), PC2 (upper right), PC3 (middle left), PC4 (middle right), PC5 (bottom left), PC6 (bottom right)

The estimates obtained from Model 1 can now be transformed back into a shape object, in order to provide a more easily interpretable result of the model in comparison to the raw estimates. This is achieved for each time-point by multiplying the appropriate predicted value in the model by its corresponding eigenvalue from the PCA. This provides a new matrix representing the change in each coordinate in comparison to the mean shape, and can simply be added to the mean shape matrix to obtain a full set of coordinates for the estimated shape object at the chosen time point. Figure 5.9 shows the estimated average face shape at each time point. It is clearly distinguishable that the biggest shape change is the size of the face, which grows bigger over time. Although other changes are very subtle, we can determine a broadening of the nose, a change in the angle of the upper lip and a change in the shape of the philtrum, which backs up the suggested interpretations of the second, fourth and fifth principal component estimates.

5.3.2 Including Sex in the Model

In order to determine any difference between the sexes, a second model (Model 2) was fitted. This model followed the same structure as Model 1, but additionally including sex as a fixed effect. The model formula is given by:

$$y_{ijkl} = \mu_k + \alpha_j + \beta_k + \gamma_l + (\alpha\beta)_{jk} + (\alpha\gamma)_{jl} + (\beta\gamma)_{kl} + (\alpha\beta\gamma)_{jkl}\delta_{ik}$$

where i denotes the subject, j the time point, k the principal component and l the sex of the subject. μ_k represents the fitted mean value, α_j the additive effect of time, β_k the additive effect of each principal component and γ_l the additive effect of sex. The model includes two-way interactions between each pair of covariates, and a three-way effect between all the covariates. Finally, δ_{ik} is the random effects vector allowing each individual to have their own principal component score value at each time point.

Figure 5.10 summarises the results of this model with the estimated val-

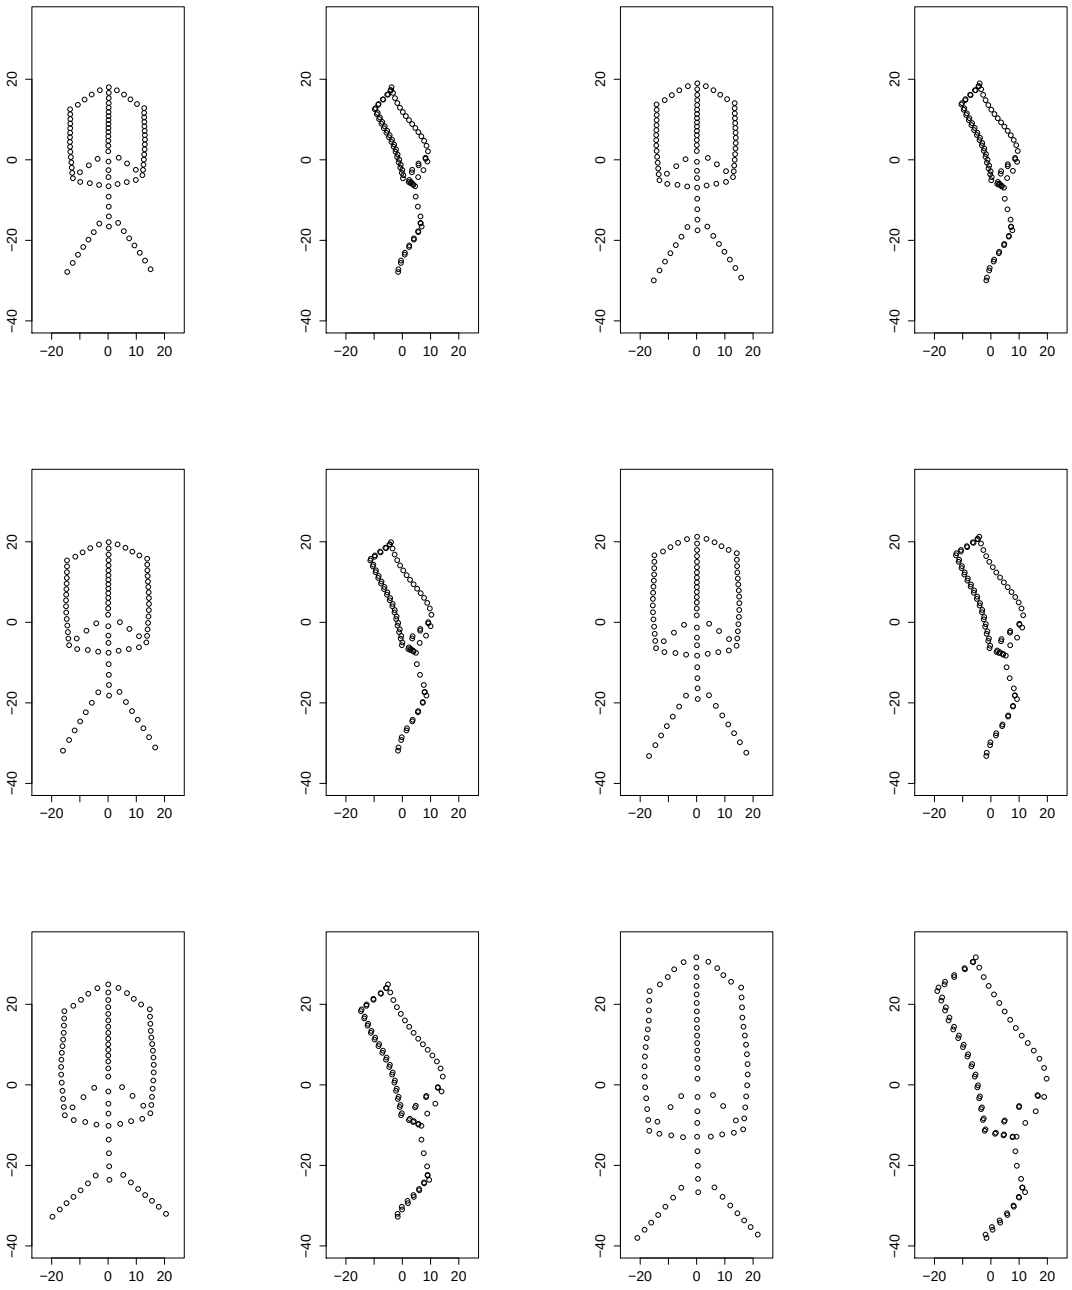

Figure 5.9: Frontal and side views of the average face shape over time according to the Model 1 estimates: 3 months (top left), 6 months (top right), 12 months (middle left), 24 months (middle right), 60 months (bottom left) and 144 months (bottom right)

ues of the fixed effects marked on the boxplots of the principal component scores for males and females (see Table D.2 in Appendix D for the full model results). Overall, the plots do not show a substantial difference between the model estimates for male and female subjects. The model estimates for the first principal component are slightly higher for males than the overall estimates for males, whereas the model estimates are slightly lower than the overall estimates for females. This tells us that males have marginally larger faces than females. However, transforming the model estimates back into shape objects reveals no discernable difference between the sexes over time (see Figure 5.11). Therefore, we can conclude that there is likely to be no significant difference in the growth pattern of males and females between the ages of 3 and 144 months in our data set.

5.3.3 Removing Size in the Procrustes Matching

As discussed in detail previously, the first principal component in the initial analysis dominated the results as the Procrustes matching had allowed size to remain. The first principal component then explained the size change in the face over time - an extremely strong effect. This led to difficulty in interpreting the latter principal components, and hence other types shape changes were masked. In order to explore whether there are any interesting shape changes, aside from size, the Procrustes matching process was repeated with scaling included, and the results of this were used to carry out the same analysis procedure described in Sections 5.3.1 and 5.3.2. The cumulative proportion of variation explained by these new principal components can be seen in Figure 5.4. Again, the first 6 principal components are required to capture over 80% of the variation in the data (see Table 5.3).

However, the first principal component still essentially represents a change in the size of the face. This can be seen in the plots of the extremes, which shows a growth in the size of the nose that is particularly noticeable in the side view, and also in the correlation observed in the plot of the centroid

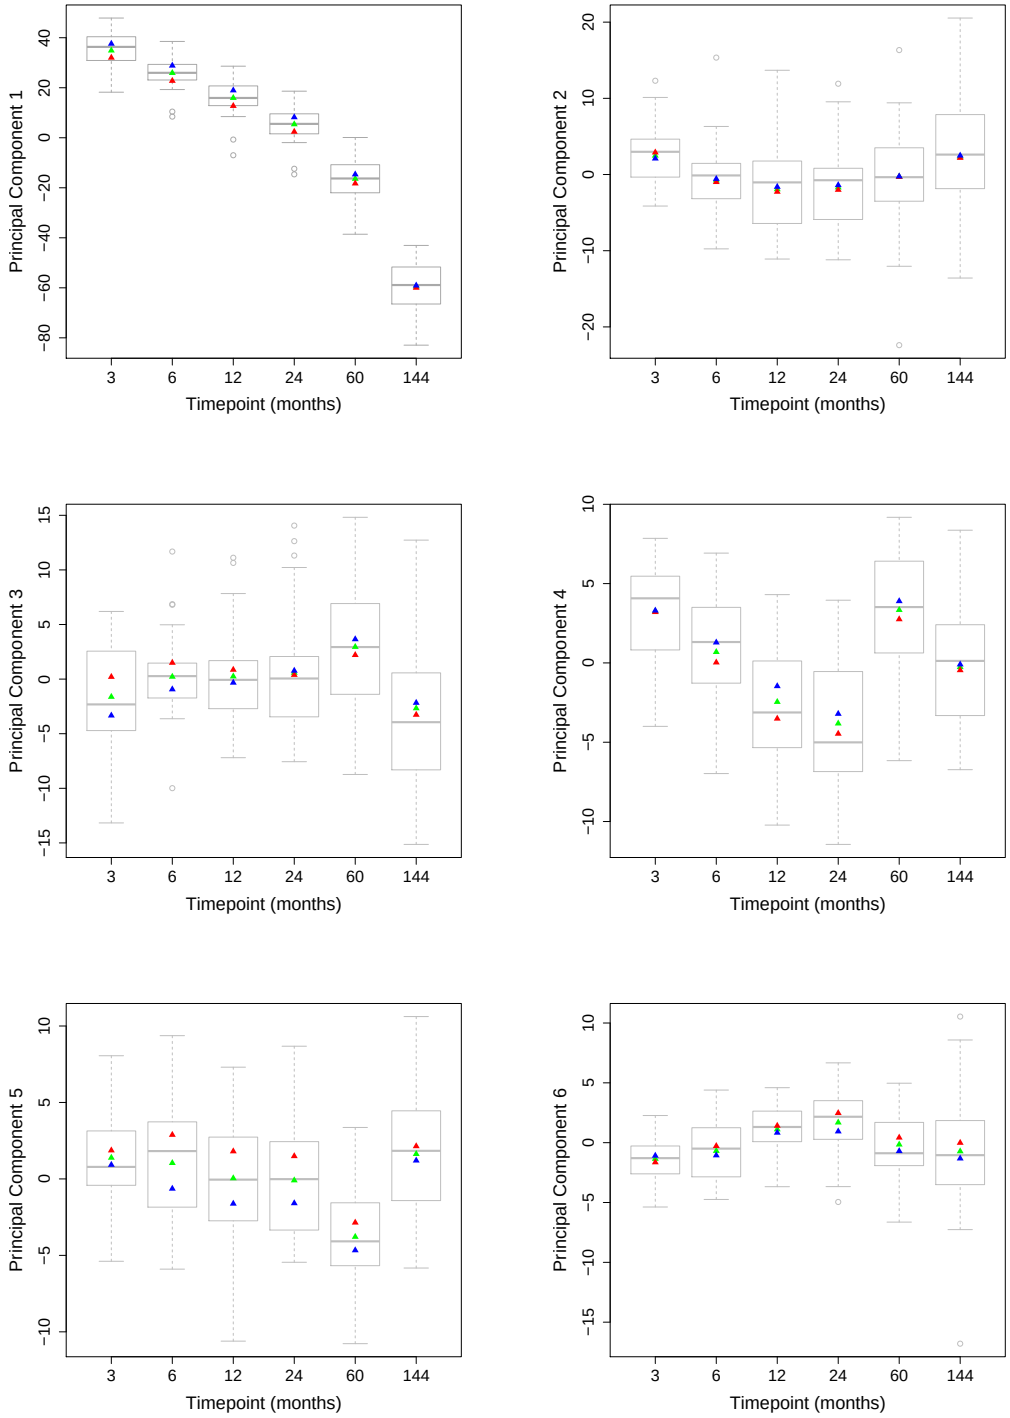

Figure 5.10: Estimates from Model 2 (blue for males and red for males) and original estimates from Model 1 (green): PC1 (upper left), PC2 (upper right), PC3 (middle left), PC4 (middle right), PC5 (bottom left), PC6 (bottom right)

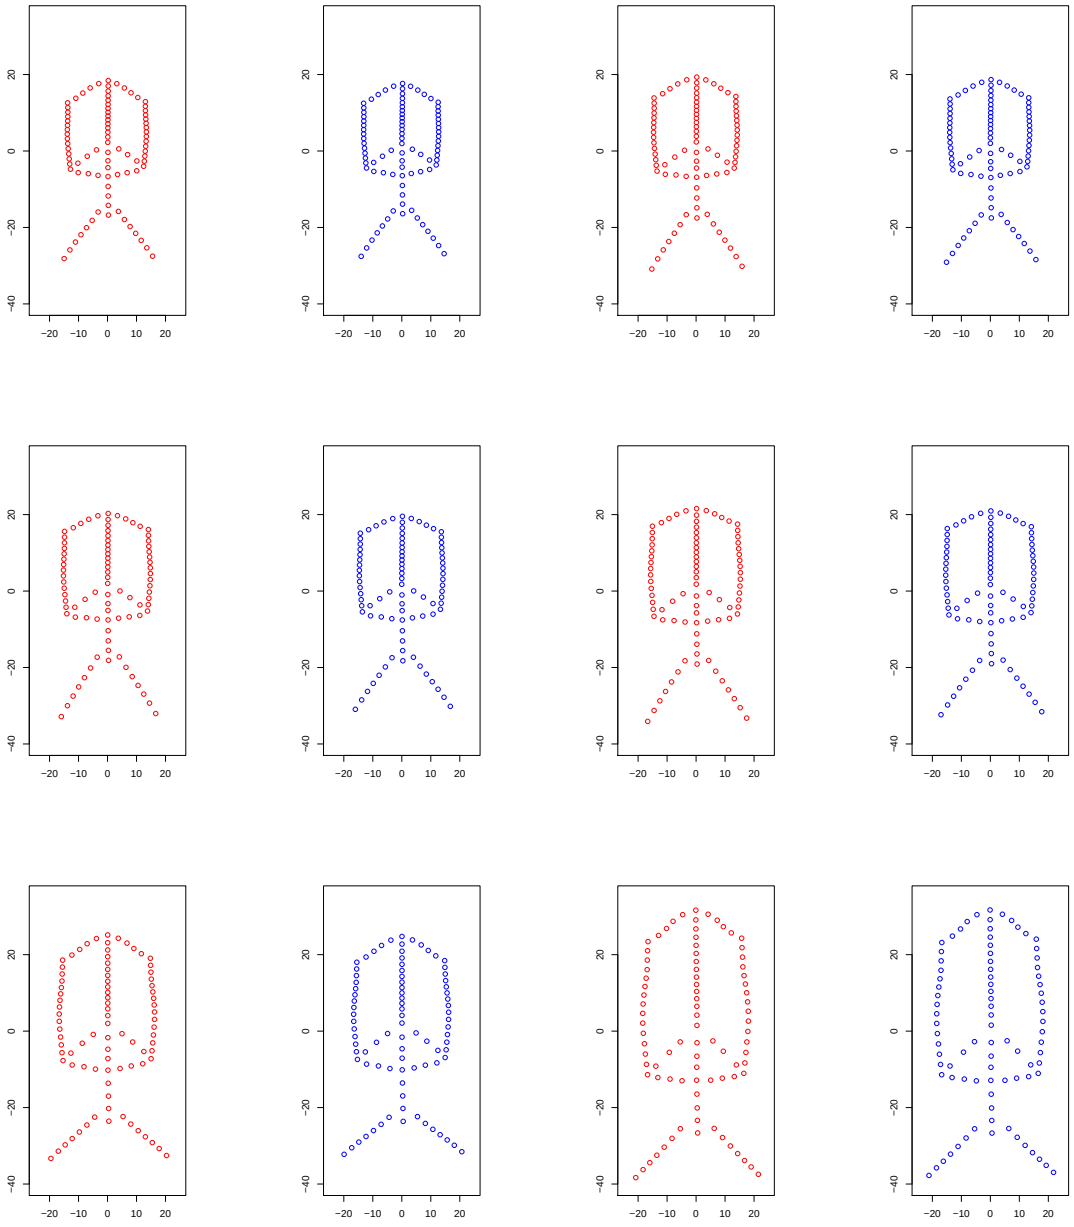

Figure 5.11: Frontal view of the average face shape over time according to the Model 2 estimates with males in blue and females in red: 3 months (top left), 6 months (top right), 12 months (middle left), 24 months (middle right), 60 months (bottom left) and 144 months (bottom right)

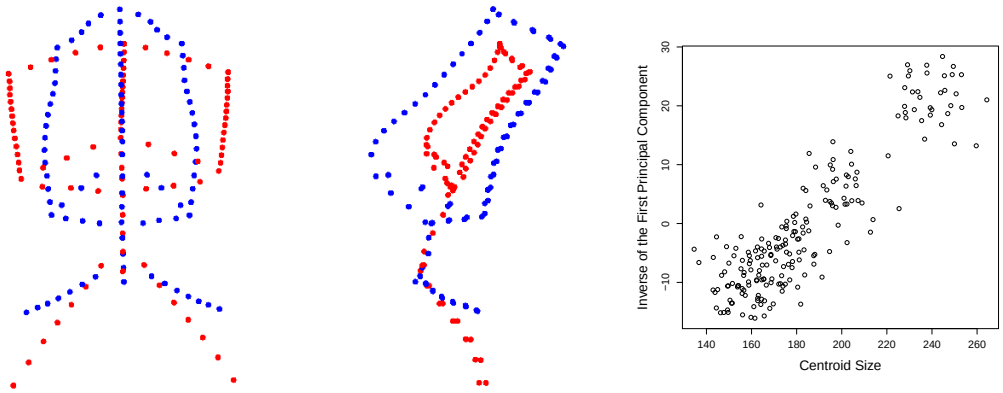

Figure 5.12: Frontal (left) and side (middle) views of the extreme values of PC1 (with scaling removed). Blue: maximum. Red: minimum. Centroid size against PC1 (right).

size against the first principal component (Figure 5.12). There is greater variability in the data in the centroid plot, compared to that in Figure 5.6, but the strong increasing trend is still very evident. This result may intuitively seem wrong at first inspection, since we believe that all size effects have been removed. However, the analysis here does not account for time since all the shapes are considered at the same time, regardless of the time point at which the data was collected. Time, size and shape are inextricably linked - over time the size of a face changes, and over time the shape of a face also changes. Therefore, the relationship we see in the centroid size plot in Figure 5.12 not only represents the link between shape change and size, but also highlights a relationship between shape change and time. Considering this, the results observed appear more reasonable. For example, if particular features of a subject's face change significantly in size, *in relation to their own face size*, then this effect will remain even once the difference in shape size *between subjects* has been removed. Hence, we have discovered that the key shape change in the faces is indeed a growth in size over time, even when the difference between shape sizes has been accounted for.

One noticeable effect is the shape change that is observed in the upper

lips in the plots of the extremes. This pattern is also observed in some of the other components (see Figure 5.13 for the full set of shapes) and was observed in some of the original component plots in Figure 5.7. The change observed is a shallowing of the angle of the upper lip, which could be related to the problem which arose with many of the infant images where the subjects mouths were open. This would stretch the upper lip and create a steeper curve. As the children age, they are able to follow instructions and with their mouths closed the upper lip angle is much smaller. A third model (Model 3) was fitted to the 6 chosen principal components using the 6 time points and allowing for a different score value for each subject at each time point. The results from this model, shown in Figure 5.15, confirm the results found previously and show that the main shape change over time is the size of the features of the face. See Table D.3 in Appendix D for the full model results.

A final model (Model 4) was fitted, including sex as a fixed effect and using the scaled principal components. Again the results of these models mirrored the previous results and indicated that no significant difference was found in the growth pattern of males and females in our cohort. Plots of the model estimates from Model 4 are shown in 5.16 and a representation of the average face shape change can be seen in 5.17. See Table D.4 in Appendix D for the full model results.

5.4 Summary

Four linear mixed models were considered in this chapter. Models 1 and 2 have scaling removed at the Procrustes matching stage. Model 1 did not include sex as a factor and Model 2 did. Models 3 and 4 allowed for scaling at the Procrustes matching stage. Model 3 did not include sex as a factor and Model 4 did. The results of the longitudinal analysis conclusively show that the growth in facial size is the overwhelmingly dominant shape change. No significant difference was found between the sexes in either Model 2 or

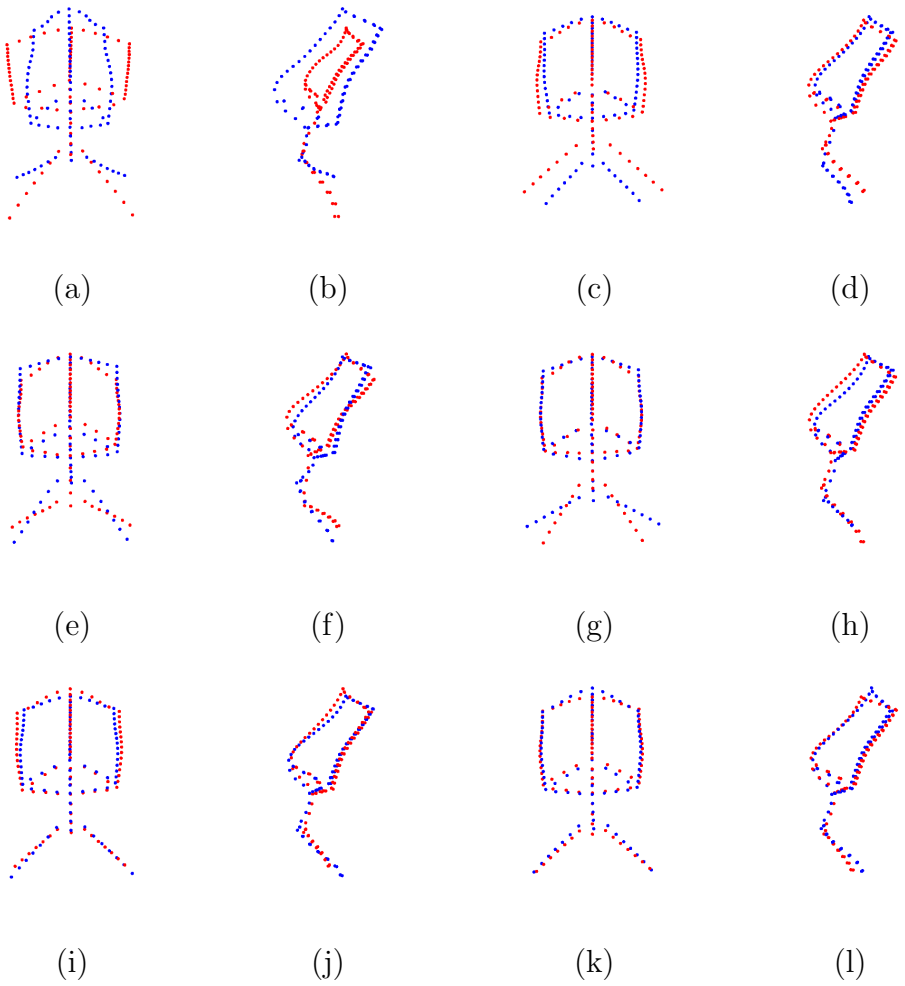

Figure 5.13: Frontal and side views of the extreme values of the principal components analysis (with scaling removed): PC1 ((a) and (b)), PC2 ((c) and (d)), PC3 ((e) and (f)), PC4 ((g) and (h)), PC5 ((i) and (j)) and PC6 ((k) and (l)). Blue: maximum. Red: minimum.

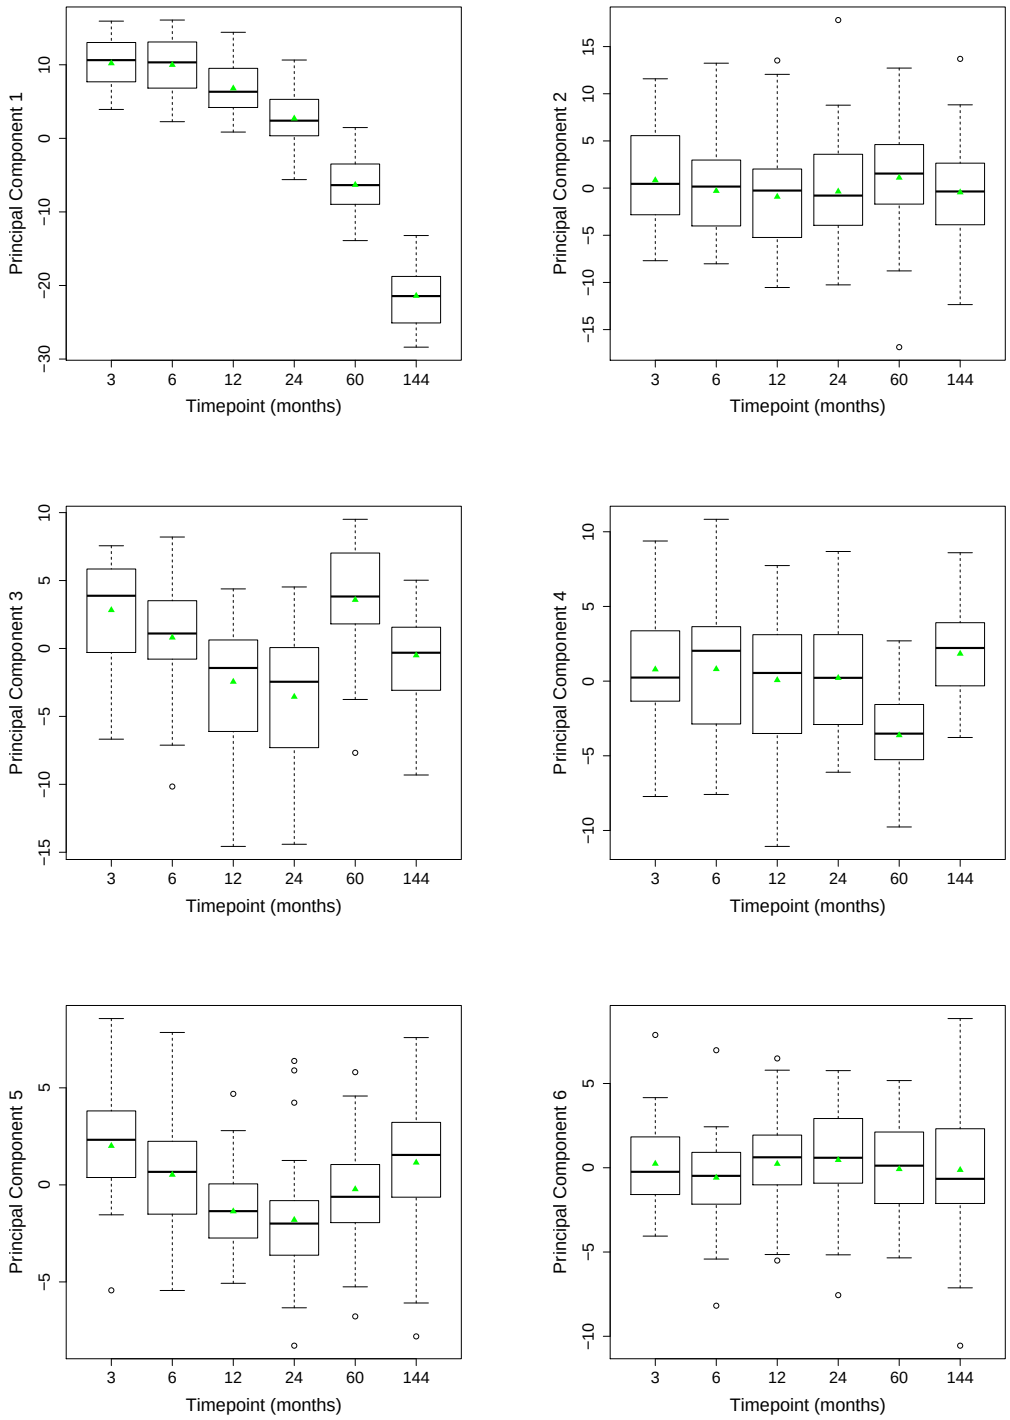

Figure 5.14: Boxplots of the first six principal components (with scaling removed) against time with estimates from Model 3 marked: PC1 (upper left), PC2 (upper right), PC3 (middle left), PC4 (middle right), PC5 (bottom left), PC6 (bottom right)

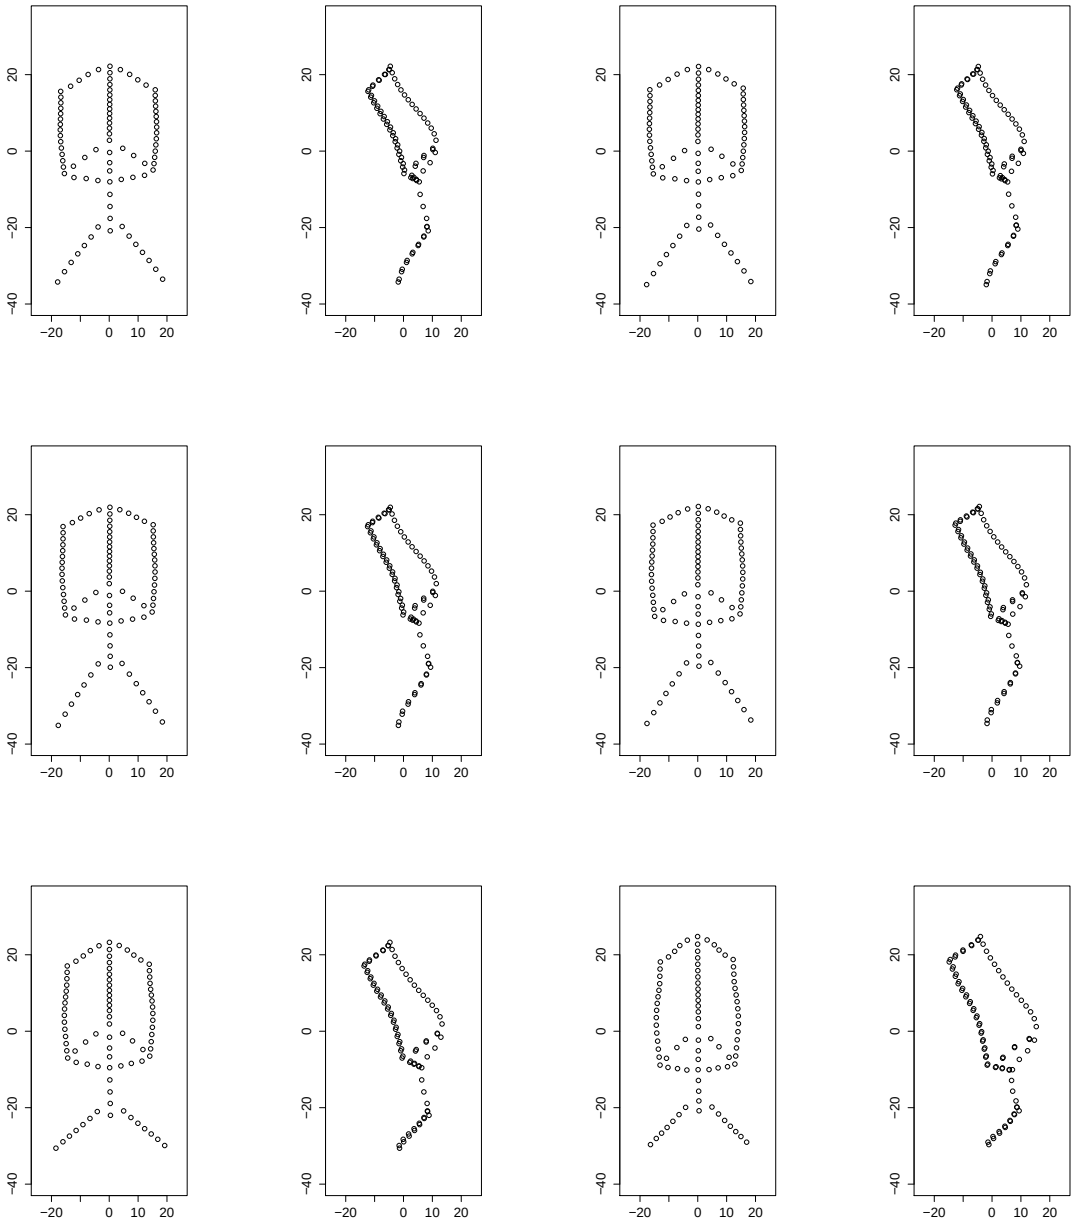

Figure 5.15: Frontal and side views of the average face shape over time according to the Model 3 estimates: 3 months (top left), 6 months (top right), 12 months (middle left), 24 months (middle right), 60 months (bottom left) and 144 months (bottom right)

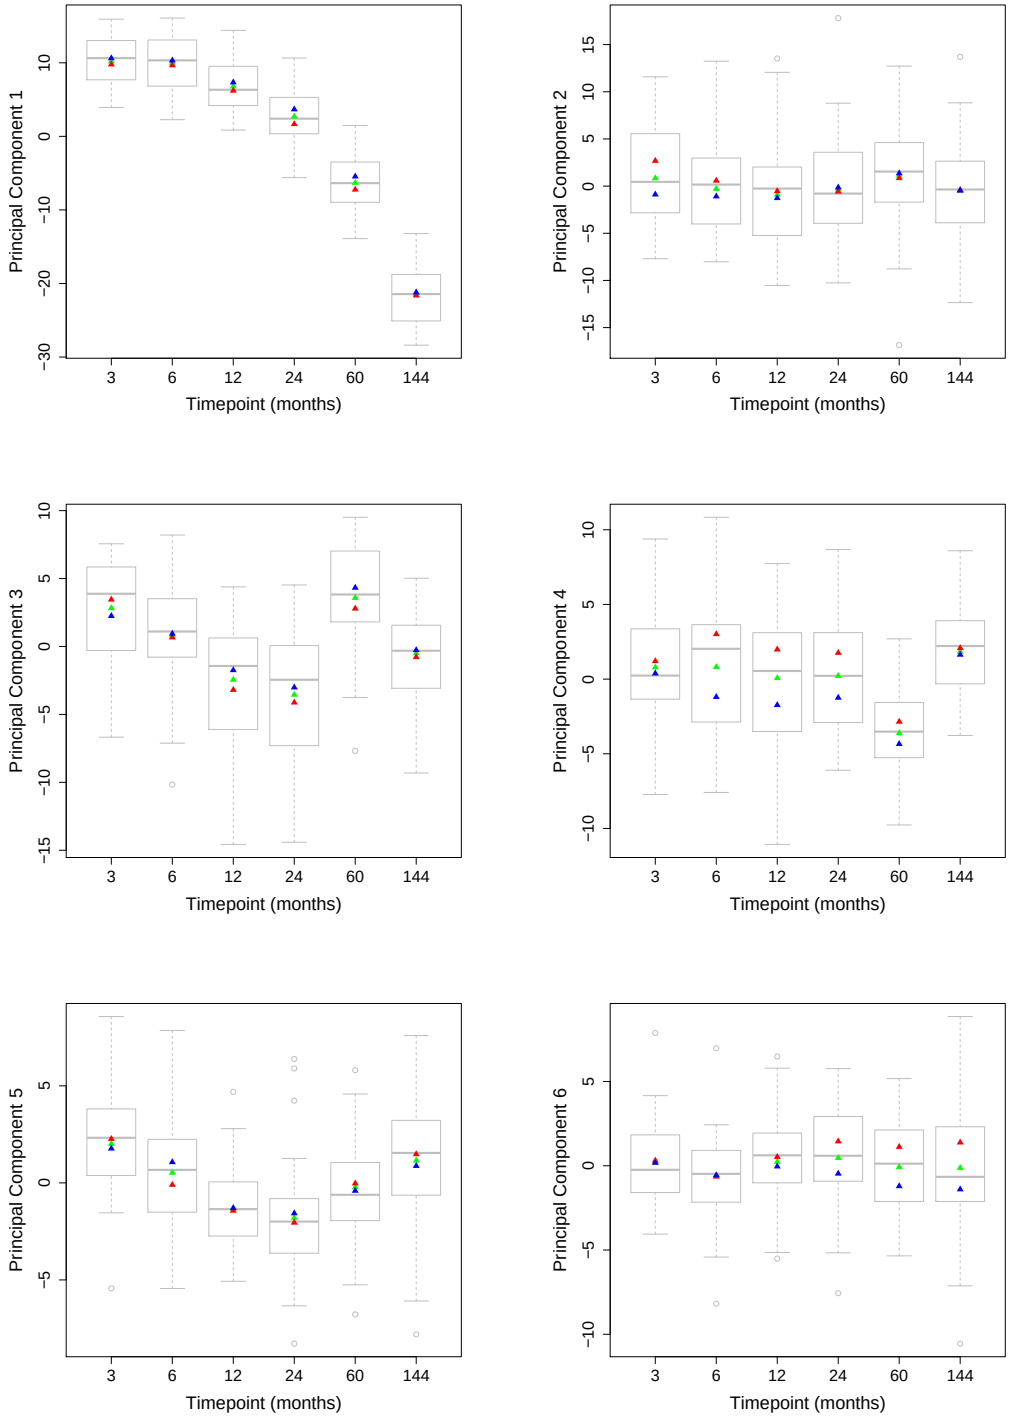

Figure 5.16: Estimates from Model 4 (blue for males and red for males) and original estimates from Model 1 (green): PC1 (upper left), PC2 (upper right), PC3 (middle left), PC4 (middle right), PC5 (bottom left), PC6 (bottom right)

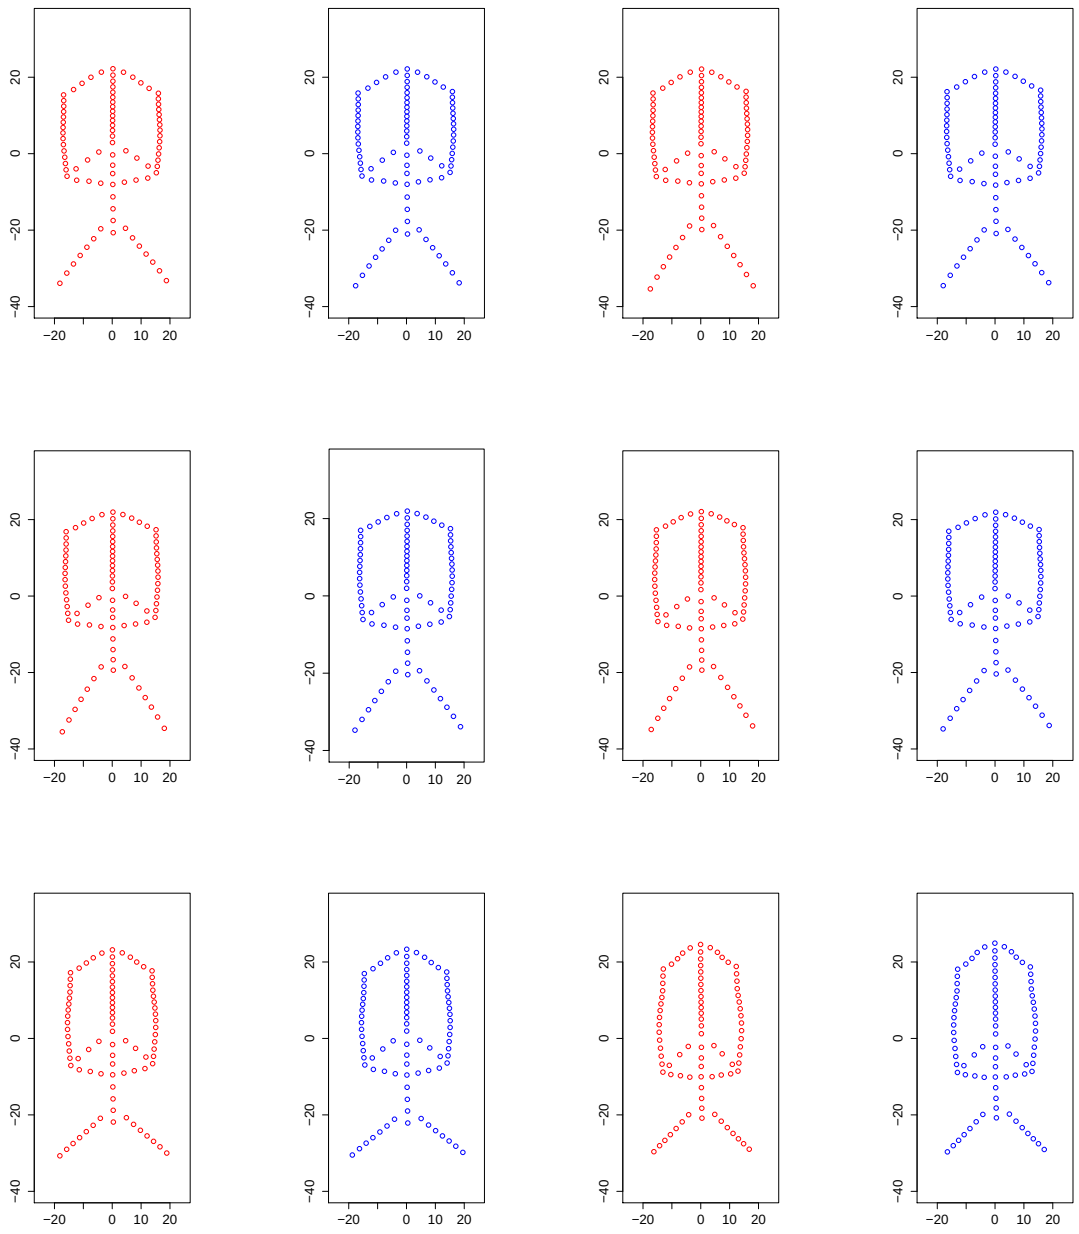

Figure 5.17: Frontal view of the average face shape over time according to the Model 4 estimates with males in blue and females in red: 3 months (top left), 6 months (top right), 12 months (middle left), 24 months (middle right), 60 months (bottom left) and 144 months (bottom right)

Model 4.

Chapter 6

Discussion and Conclusions

6.1 Summary

The key objectives of this thesis were to obtain a database of images of a cohort of 12-year old children who participated in a longitudinal study as infants, to develop methods of automatic curve identification and to analyse the longitudinal data using appropriate curve and statistical shape analyses. This thesis covered the initial capture of a control data set of facial images of 100 children, who were originally recruited at the age of 3 months and followed-up to 5 years. At the 12-year time point 39 subjects were successfully followed-up, although, due to some issues with the data that was collected, 4 cases had to be excluded from statistical analysis.

Methods of curve identification were discussed, specifically for the curves of the upper and lower lips and the midline of the lips. An automatic curve identification algorithm using the shape index was developed, and compared to planepath curves.

The landmark configurations were identified on the facial images collected at 12 years, and these, along with the landmark configurations from the previous 5 time points, were analysed in depth in Chapter 5. The results of the longitudinal analysis conclusively show that the growth in facial size is the prominent shape change. The width, length and protrusion of the nose

increases over time, and the length of the philtrum and upper lip grows. The effect of size is so strong that it potentially masks other, more subtle, shape changes. No significant difference was found between the sexes in any of our analysis.

6.2 Limitations to the Data

One of the issues in the data collection process described in Section 2.2 was the lack of up-to-date contact details for the subjects' families. This meant that many of the subjects had to be excluded completely from the study as there was no information available, and even in the cases where information was available, some parents did not respond to any of the various methods of communication attempted. Despite two letters being sent to the last known postal address of each family, a letter being sent to a secondary contact and an email being sent where addresses were available, there were only 39 responses of a potential 59. Since 88 subjects participated up until the 5 year time point, it would have been desirable to obtain a follow-up rate as close to this number as possible. However, there was only contact information on 59 of these 88 subjects, and in the end 39 subjects returned to have their images were captured.

A problem that was encountered when the imaging sessions began was that some of the subjects had fixed braces, common for the current age of the children. This could have the potential to affect the natural lip shape, and increase the protrusion of the upper or lower lips in particular. Unfortunately there was no way to avoid an image with the fixed braces, as they are not removable. If the subject had fixed braces, this was noted in the database detailing information on the subjects. In future, more detailed, analysis this information could be utilised in order to determine whether the braces do have a significant effect on the natural shape of the mouth.

Four of the images captured during the follow-up process had to be ex-

cluded from the study. Two of these were due to the lips being parted slightly and two were due to a distortion of the three-dimensional image after the image building process. The latter two images are probably not useable in any future analysis as the distortion occurs around some of the key features of the face, such as the nose or mouth. However, it would potentially be possible to include the former two images if a different landmarking system was used, assigning two landmarks to the centre of the mouth, one at the bottom of the upper lip and one at the top of the lower lip (in cases where the lips are not parted, a double "dummy" landmark would need to be added). Ideally the images should be captured with the mouths completely closed and so closer inspection of the captured images during the session may avoid this situation.

This problem of a subject's mouth being open is more of a concern in the infant data, between the 3 and 24 month time points. As an infant may not understand or be willing to comply with an instruction, this issue arose in a large proportion of cases. In a few cases multiple images were captured in order to achieve the correct facial expression and so the image with the open mouth could be discarded in favour of the one with the mouth closed. However, for other subjects multiple images were not taken and in total 110 of 133 subjects between the ages of 3 and 24 months had their mouths open. In order to accurately represent the face shape and allow for comparisons between subjects and between time points, it was decided that only the region above (and including) the upper lip would be considered in the longitudinal analysis. Clearly this discards vital information which could be useful in the analysis of facial shape change. Additionally, the results in Chapter 5 indicate that the open mouths may have had an influence on the shape change over time. By the 5 and 12 year time points, all of the subjects were able to follow the instructions to keep their mouths closed. One of the shape changes observed (particularly once scale had been removed from the analysis) was a shallowing of the angle of the upper lip. Having the mouth

open stretches the upper lip and creates a steeper curve, whereas when the mouth is at rest and closed the upper lip angle is much smaller. This issue is, unfortunately, mostly unavoidable when dealing with such young infants. In future studies it may be possible to achieve the desired facial expression if many images are captured, but this could require imaging sessions which are, in practise, too long.

6.3 Limitations to the Curve Identification Methods

In Chapter 4 two methods of curve identification were explored. The first of these methods made use of the shape index and principal curves in order to fit smooth curves between two chosen landmarks. The lips were the focus of the analysis. The current form of this algorithm is able to produce a flexible curve which captures some subtle features of the lip curves. However, the curves do not necessarily lie on the facial surface and therefore are not always as precise as would be desired. If this algorithm could be developed further, then it may be applicable to other areas of the face. This may be especially useful in the regions of the face that are notoriously hard to capture, such as the brow and jaw lines.

The second method considered in Chapter 4 was the planepath curve, which fits a curve that describes the shortest distance along the facial surface between two points of interest. The planepath curves were considered to be advantageous for use in the analysis as they sit on the surface of the shape and can be accurately applied to regions of the face such as the nose. However, it is not currently able to handle more difficult curves such as the brow and jaw lines. This means that the analysis in this thesis has been limited to the central features of the face. The results in Chapter 5 found no difference in facial shape change between males and females from 3 to 144 months. This may seem like a surprising result, but since most of the

subject have not yet reached puberty it is plausible that the distinguishing features between the sexes are not yet apparent. However, it has been suggested that a distinguishing feature between males and females may in fact lie in the shape of the brow and jaw line (reference?). Since we do not have this information available, this hypothesis cannot be explored. Once curve identification methods are able to accurately represent these types of facial curves, similar analysis could be repeated including this extra information.

6.4 Areas for Further Study

It would be advantageous to repeat the analysis carried out in Chapter 5 with a larger set of curves, including the midline, lower lip and brow and jaw lines. This would provide a more comprehensive representation of the face and may highlight new results that were not possible with the current set of curves.

It would also be of interest to follow-up the subjects in the longitudinal study at multiple stages over the next few years. As the children progress through puberty there may be interesting results, particularly in the difference between the sexes as discussed above. At the end of each image capturing session the possibility of follow-ups was discussed with the subjects and their parents. All participants indicated that they would be happy to continue their participation in the study.

Appendix A

Approach Letter

University of Glasgow | School of Mathematics & Statistics

From: Prof. Adrian Bowman
Professor of Statistics
Tel: 0141-330-4046
E-mail: adrian.bowman@glasgow.ac.uk

November 2, 2012

Dear

A study of facial shape

Several years ago, you and your child were kind enough to participate in a study of facial shape in children which was conducted by a team from the University of Glasgow. The project was led by Prof. Ayoub and the data were collected by Dr. Jill White, both from Glasgow Dental School. Although we did not meet, I was also a member of the team, responsible for analysis of the images obtained. The data collected from your child and from others provided invaluable information on normal facial shape and in particular it allowed comparisons to be made with the facial shape of children who were born with a cleft lip and had undergone corrective surgery. The results of the study were highly rated by the Chief Scientists's Office, who had provided funding, and they were published in the scientific literature. Thank you for your assistance.

I am contacting you again because there is now an opportunity to follow up the study, to look at facial shape in older children. It would be possible to recruit new cases to study this at particular ages but there are very great advantages in continuing to collect data from the same group of children whose images were taken in earlier years. This would allow us to track the changes in individual children and so it would be much more informative about how growth affects shape. Collecting images from the same group of children throughout childhood would create an invaluable dataset. I am therefore writing to ask whether you and your child would be willing to visit us for a further three-dimensional photograph to be taken.

The practical arrangements would involve a trip to the Mathematics & Statistics building in University Gardens, on the main campus of the University. We need a parent or guardian to be present, but the day and time can be arranged to suit you. The photograph should be taken within 10 minutes or so. Anna Price, a research student in the University, would contact you to discuss details and it is Anna who would take the photograph. We do not have the funds to cover travel costs but if you are able to participate then we would be delighted to explain the nature of the research which is going on, if you are interested. We would also provide you and your child with a three-dimensional facial photograph, which you might find fun to have.

I should reassure you that the project has been approved by the Ethics Committee of the College of Science and Engineering at the University of Glasgow. In particular, any data you provide will be held securely within the University of Glasgow. Your child's image will not be used in any publications (unless we approach you explicitly about that and you give us specific permission). You are free at any stage, even after any image has been taken, to request that your child is withdrawn from the study and that any data we hold is deleted.

We are also collecting facial shape from adults so if you would like to have your own photograph taken we would be delighted to do that. An information sheet about adult data collection is attached.

If you have any questions about this we would be very happy to answer them. If you would like to contact Anna (a.price.1@research.gla.ac.uk; phone number - *to be added*), she would be delighted to explain the background and to make arrangements for a visit. You are also very welcome to contact me directly (adrian.bowman@glasgow.ac.uk; telephone 0141-330-4046). If we don't hear from you within a couple of weeks, I hope you won't mind if we try to contact you again, in case we have an out-of-date address.

Yours sincerely,

Prof. Adrian Bowman

Appendix B

Permission Sheet

University of Glasgow | School of Mathematics & Statistics

Research Project: A Study on Facial Shape

I have read the letter about this project and I agree that the data collected from the child identified below can be used for the purposes of the research involved.

I understand that I am free to withdraw my child from this study at any stage and to request that the associated data be deleted.

Study Number……..

Name……..

Name of Parent/Guardian……

Signature of Parent/Guardian……

Signature of experimenter……

Date……

Appendix C

Data Capture Protocol

Data Capture Protocol

This document is intended to outline a standard format for obtaining information from participants in a 3D image capturing project.

Questionnaire

All participants will be asked a series of questions, as documented in the participant information sheet. The analyst will ask these questions and input the answers directly into a secure database.

Equipment

Throughout a capture session with multiple people, the camera equipment will be positioned in the same location, or as close to it as possible.

The camera will be calibrated at the beginning of each day and again throughout the day as required, at the discretion of the analyst.

A static chair will be positioned in front of the camera. The cameras will be positioned so that the participant's eyes are lined up at the centre of each camera view before the image is taken.

It is expected that three images will be captured and from these, the best set will be selected. The selected image will then be built into a 3D image and then checked to ensure that the image is useable for the study.

Participants

Each of the participants will be allocated a set time in which to arrive at the location of the data collection. Due to the nature of some of the questions that will be asked during the questionnaire, only one participant will be in the room along with the analyst at one time.

Each subject will be seated in the static chair. All participants will be asked to pull their hair away from their face using hair clips or hair bands provided. Prior to their arrival participants will also be asked to refrain from using heavy make-up and to shave any facial hair. A shiny or hairy facial surface could produce an image where the facial structure or boundaries of facial features that are inaccurate. Facial jewellery should be removed, if that is easy to do, but embedded jewellery such as studs should not be removed in view of the complications which may be involved. Assessment of image quality can be made after the capture.

Participants will be asked to tilt their heads slightly backwards so that the area around the nose and chin can be captured accurately. The participant's eyes should be open. Participants will be asked to say "Mississippi", followed by the letter "n" and then asked to let their mouth rest, in order to relax the mouth and jaw into a natural position before capturing the image. If a participant is clenching their jaw or holding their mouth in an unnatural position, this will give an inaccurate measurement of their true facial structure.

Where rest position involves an open mouth, participants will additionally be asked to close their lips lightly together and a further image will be captuted.

Timing

It is estimated that collection of information will take approximately 15 minutes per participant.

Appendix D

Additional Results from Chapter 5

Table D.1: Summary of Fixed and Random Effects in Model 1

Fixed Effects			
Component	Estimate	Standard Error	t-value
Intercept	34.92	0.98	35.64
Time point 6	-9.00	0.78	-11.48
Time point 12	-19.00	0.78	-24.43
Time point 24	-29.53	0.78	-37.98
Time point 60	-51.33	0.78	-66.02
Time point 144	-94.41	0.78	-121.43
PC 2	-32.45	1.46	-22.21
PC 3	-36.53	1.45	-25.14
PC 4	-31.68	1.23	-25.77
PC 5	-33.54	1.32	-25.37
PC 6	-36.27	1.17	-30.88
Time point 6 * PC 2	5.74	21.10	5.18
Time point 12 * PC 2	14.59	1.10	13.27
Time point 24 * PC 2	25.37	1.10	23.08
Time point 60 * PC 2	48.58	1.10	44.20
Time point 144 * PC 2	94.28	1.10	85.79
Time point 6 * PC 3	10.81	1.11	9.75
Time point 12 * PC 3	20.85	1.10	18.95
Time point 24 * PC 3	31.72	1.10	28.83
Time point 60 * PC 3	55.89	1.10	50.80
Time point 144 * PC 3	93.35	1.10	84.85
Time point 6 * PC 3	6.44	1.11	5.82
Time point 12 * PC 4	13.29	1.10	12.10
Time point 24 * PC 4	22.47	1.10	20.46
Time point 60 * PC 4	51.43	1.10	46.83
Time point 144 * PC 4	90.91	1.10	82.77
Time point 6 * PC 5	8.65	1.11	7.81
Time point 12 * PC 5	17.66	1.10	16.06
Time point 24 * PC 5	28.06	1.10	25.52
Time point 60 * PC 5	46.17	1.10	41.98
Time point 144 * PC 5	94.66	1.10	86.08
Time point 6 * PC 6	9.64	1.11	8.71
Time point 12 * PC 6	21.45	1.10	19.53
Time point 24 * PC 6	32.56	1.10	29.65
Time point 60 * PC 6	52.52	1.10	47.82
Time point 144 * PC 6	95.03	1.10	86.52

Random Effects	
Component	Standard Deviation
Intercept	4.68
PC 2	7.16
PC 3	7.10
PC 4	5.44
PC 5	6.14
PC 6	4.99

Table D.2: Summary of Fixed and Random Effects in Model 2

Fixed Effects			
Component	Estimate	Standard Error	t-value
Intercept	37.48	1.24	30.13
Time Point 6	-8.81	1.07	-8.25
Time Point 12	-18.71	1.07	-17.52
Time Point 24	-29.45	1.07	-27.59
Time Point 60	-52.31	1.07	-48.99
Time Point 144	-96.71	1.07	-90.76
PC 2	-35.59	1.68	-21.21
PC 3	-40.82	1.88	-21.72
PC 4	-34.23	1.73	-19.77
PC 5	-36.56	1.60	-22.94
PC 6	-38.60	1.55	-24.98
Sex (male)	-5.27	1.62	-3.25
Time Point 6 *PC 2	6.13	1.51	4.06
Time Point 12 *PC 2	14.99	1.51	9.93
Time Point 24 *PC 2	25.96	1.51	17.20
Time Point 60 *PC 2	49.96	1.51	33.11
Time Point 144 *PC 2	97.11	1.50	64.55
Time Point 6 *PC 3	11.16	1.51	7.39
Time Point 12 *PC 3	21.67	1.51	14.35
Time Point 24 *PC 3	33.51	1.51	22.18
Time Point 60 *PC 3	59.25	1.51	39.23
Time Point 144 *PC 3	97.79	1.51	64.80
Time Point 6 *PC 4	6.77	1.51	4.49
Time Point 12 *PC 4	13.92	1.51	9.23
Time Point 24 *PC 4	22.92	1.51	15.20
Time Point 60 *PC 4	52.88	1.51	35.06
Time Point 144 *PC 4	93.28	1.50	62.06
Time Point 6 *PC 5	7.24	1.51	4.80
Time Point 12 *PC 5	16.16	1.51	10.71
Time Point 24 *PC 5	26.94	1.51	17.85
Time Point 60 *PC 5	46.72	1.51	30.95
Time Point 144 *PC 5	96.99	1.50	64.48
Time Point 6 *PC 6	8.78	1.51	5.82
Time Point 12 *PC 6	20.57	1.51	13.64
Time Point 24 *PC 6	31.43	1.51	20.84
Time Point 60 *PC 6	52.64	1.51	34.91
Time Point 144 *PC 6	96.38	1.50	64.21
Time Point 6 *Sex (male)	-0.45	1.55	-0.29
Time Point 12 *Sex (male)	-0.59	1.54	-0.38
Time Point 24 *Sex (male)	-0.16	1.54	-0.10
Time Point 60 *Sex (male)	2.02	1.54	1.31
Time Point 144 *Sex (male)	4.70	1.56	3.00
PC 2 *Sex (male)	6.50	2.20	2.95
PC 3 *Sex (male)	6.49	2.41	3.66
PC 4 *Sex (male)	5.29	2.23	2.38
PC 5 *Sex (male)	6.23	2.14	2.91
PC 6 *Sex (male)	4.82	2.09	2.31

Component	Estimate	Standard Error	t-value
Time Point 6 * PC 2 * Sex (male)	-0.75	2.19	-0.34
Time Point 12 * PC 2 * Sex (male)	0.85	2.17	-0.39
Time Point 24 * PC 2 * Sex (male)	-1.25	2.17	-0.58
Time Point 60 * PC 2 * Sex (male)	-2.87	2.17	-1.32
Time Point 144 * PC 2 * Sex (male)	-5.81	2.21	-2.63
Time Point 6 * PC 3 * Sex (male)	-0.61	2.19	-0.28
Time Point 12 * PC 3 * Sex (male)	-1.71	2.17	-0.79
Time Point 24 * PC 3 * Sex (male)	-3.70	2.17	-1.70
Time Point 60 * PC 3 * Sex (male)	-6.94	2.17	-3.19
Time Point 144 * PC 3 * Sex (male)	-9.19	2.22	-4.15
Time Point 6 * PC 4 * Sex (male)	-0.70	2.19	-0.32
Time Point 12 * PC 4 * Sex (male)	-0.69	2.17	-0.61
Time Point 24 * PC 4 * Sex (male)	-0.97	2.17	-0.45
Time Point 60 * PC 4 * Sex (male)	-3.03	2.17	-1.40
Time Point 144 * PC 4 * Sex (male)	-4.89	2.20	-2.22
Time Point 6 * PC 5 * Sex (male)	3.02	2.19	1.38
Time Point 12 * PC 5 * Sex (male)	3.06	2.17	1.41
Time Point 24 * PC 5 * Sex (male)	2.29	2.17	1.06
Time Point 60 * PC 5 * Sex (male)	-1.15	2.17	-0.53
Time Point 144 * PC 5 * Sex (male)	-4.73	2.21	-2.14
Time Point 6 * PC 6 * Sex (male)	1.84	2.19	0.84
Time Point 12 * PC 6 * Sex (male)	1.80	2.17	0.83
Time Point 24 * PC 6 * Sex (male)	2.33	2.17	1.07
Time Point 60 * PC 6 * Sex (male)	-0.26	2.17	-0.12
Time Point 144* PC 6 * Sex (male)	-2.66	2.20	-1.21

Random Effects	
Component	Standard Deviation
Intercept	4.61
PC 2	5.99
PC 3	7.23
PC 4	6.44
PC 5	5.33
PC 6	5.01

Table D.3: Summary of Fixed and Random Effects in Model 3

Fixed Effects			
Component	Estimate	Standard Error	t-value
Intercept	10.21	0.65	15.64
Time Point 6	-0.22	0.68	-0.33
Time Point 12	-3.42	0.68	-5.05
Time Point 24	-7.51	0.68	-11.10
Time Point 60	-16.53	0.68	-24.44
Time Point 144	-31.61	0.68	-46.74
PC 2	-31.61	1.08	-8.67
PC 3	-31.68	0.99	-7.49
PC 4	-9.41	0.96	-9.85
PC 5	-8.19	0.79	-10.38
PC 6	-9.96	0.93	-10.70
Time Point 6 * PC 2	-0.91	0.96	-0.94
Time Point 12 * PC 2	1.67	0.96	1.74
Time Point 24 * PC 2	6.32	0.96	6.60
Time Point 60 * PC 2	16.80	0.96	17.55
Time Point 144 * PC 2	30.33	0.96	31.68
Time Point 6 * PC 3	-1.80	0.96	-1.87
Time Point 12 * PC 3	-1.86	0.96	-1.95
Time Point 24 * PC 3	1.13	0.96	1.18
Time Point 60 * PC 3	17.28	0.96	18.06
Time Point 144 * PC 3	28.28	0.96	29.56
Time Point 6 * PC 4	0.24	0.96	0.25
Time Point 12 * PC 4	2.69	0.96	2.81
Time Point 24 * PC 4	6.93	0.96	7.24
Time Point 60 * PC 4	12.12	0.96	12.66
Time Point 144 * PC 4	32.65	0.96	34.12
Time Point 6 * PC 5	-1.27	0.96	-1.32
Time Point 12 * PC 5	0.04	0.95	0.04
Time Point 24 * PC 5	3.69	0.95	3.87
Time Point 60 * PC 5	14.29	0.95	14.99
Time Point 144 * PC 5	30.75	0.95	32.24
Time Point 6 * PC 6	-0.61	0.96	-0.63
Time Point 12 * PC 6	3.40	0.96	3.56
Time Point 24 * PC 6	7.73	0.96	8.07
Time Point 60 * PC 6	16.20	0.96	16.93
Time Point 144 * PC 6	31.24	0.96	32.63

Random Effects	
Component	Standard Deviation
Intercept	2.460
PC 2	4.81
PC 3	4.03
PC 4	3.77
PC 5	2.07
PC 6	3.55

Table D.4: Summary of Fixed and Random Effects in Model 4

Fixed Effects			
Component	Estimate	Standard Error	t-value
Intercept	10.54	0.85	12.34
Time Point 6	-0.33	0.94	-0.36
Time Point 12	-0.31	0.94	-3.54
Time Point 24	-6.98	0.94	-7.46
Time Point 60	-16.10	0.94	-17.21
Time Point 144	-31.87	0.93	-34.25
PC 2	-11.60	1.49	-7.78
PC 3	-8.35	1.29	-6.47
PC 4	-10.24	1.26	-8.15
PC 5	-8.86	1.06	-8.33
PC 6	-10.42	1.15	-9.07
Sex (male)	-0.66	1.20	-0.55
Time Point 6 * PC 2	0.13	1.32	0.10
Time Point 12 * PC 2	2.93	1.32	2.21
Time Point 24 * PC 2	7.73	1.32	5.83
Time Point 60 * PC 2	18.34	1.32	13.85
Time Point 144 * PC 2	32.29	1.32	24.47
Time Point 6 * PC 3	-0.97	1.32	-0.74
Time Point 12 * PC 3	-0.68	1.32	-0.52
Time Point 24 * PC 3	1.71	1.32	1.30
Time Point 60 * PC 3	18.17	1.32	13.74
Time Point 144 * PC 3	29.36	1.32	22.31
Time Point 6 * PC 4	-1.24	1.32	-0.94
Time Point 12 * PC 4	1.19	1.32	0.90
Time Point 24 * PC 4	5.35	1.32	4.04
Time Point 60 * PC 4	11.38	1.32	8.59
Time Point 144 * PC 4	33.15	1.32	25.16
Time Point 6 * PC 5	-0.36	1.32	-0.27
Time Point 12 * PC 5	0.25	1.32	0.19
Time Point 24 * PC 5	3.65	1.32	2.77
Time Point 60 * PC 5	13.93	1.32	10.56
Time Point 144 * PC 5	31.00	1.31	23.70
Time Point 6 * PC 6	-0.41	1.32	-0.31
Time Point 12 * PC 6	3.09	1.32	2.33
Time Point 24 * PC 6	6.33	1.32	4.79
Time Point 60 * PC 6	14.71	1.32	11.12
Time Point 144 * PC 6	30.28	1.32	23.02
Time Point 6 * Sex (male)	0.21	1.36	0.16
Time Point 12 * Sex (male)	-0.25	1.34	-0.19
Time Point 24 * Sex (male)	-1.13	1.34	-0.84
Time Point 60 * Sex (male)	-0.91	1.34	-0.68
Time Point 144 * Sex (male)	0.49	1.36	0.36
PC 2 * Sex (male)	4.56	1.94	2.36
PC 3 * Sex (male)	1.98	1.78	1.12
PC 4 * Sex (male)	1.64	1.75	0.94
PC 5 * Sex (male)	1.33	1.52	0.87
PC 6 * Sex (male)	0.90	1.63	0.55

Component	Estimate	Standard Error	t-value
Time Point 6 * PC 2 * Sex (male)	-2.11	1.92	-1.10
Time Point 12 * PC 2 * Sex (male)	-2.11	1.90	-1.36
Time Point 24 * PC 2 * Sex (male)	-2.88	1.90	-1.51
Time Point 60 * PC 2 * Sex (male)	-3.16	1.90	-1.66
Time Point 144 * PC 2 * Sex (male)	-4.00	1.94	-2.07
Time Point 6 * PC 3 * Sex (male)	-1.70	1.92	-0.89
Time Point 12 * PC 3 * Sex (male)	-2.41	1.90	-1.27
Time Point 24 * PC 3 * Sex (male)	-1.19	1.90	-0.63
Time Point 60 * PC 3 * Sex (male)	-1.83	1.90	-0.96
Time Point 144 * PC 3 * Sex (male)	-2.22	1.93	-1.15
Time Point 6 * PC 4 * Sex (male)	3.16	1.92	1.65
Time Point 12 * PC 4 * Sex (male)	3.14	1.90	1.65
Time Point 24 * PC 4 * Sex (male)	3.30	1.90	1.73
Time Point 60 * PC 4 * Sex (male)	1.58	1.90	0.83
Time Point 144 * PC 4 * Sex (male)	0.93	1.93	-0.48
Time Point 6 * PC 5 * Sex (male)	-1.89	1.91	-0.99
Time Point 12 * PC 5 * Sex (male)	-0.39	1.90	-0.21
Time Point 24 * PC 5 * Sex (male)	0.14	1.90	0.08
Time Point 60 * PC 5 * Sex (male)	0.79	1.90	0.41
Time Point 144 * PC 5 * Sex (male)	-0.43	1.91	-0.22
Time Point 6 * PC 6 * Sex (male)	-0.41	1.92	-0.22
Time Point 12 * PC 6 * Sex (male)	0.69	1.90	0.36
Time Point 24 * PC 6 * Sex (male)	2.91	1.90	1.53
Time Point 60 * PC 6 * Sex (male)	3.12	1.90	1.64
Time Point 144 * PC 6 * Sex (male)	2.20	1.93	1.14

Random Effects	
Component	Standard Deviation
Intercept	2.24
PC 2	5.43
PC 3	3.86
PC 4	3.56
PC 5	1.86
PC 6	2.62

Bibliography

Ayoub, A., Garrahy, A., Hood, C., White, J., Bock, M., Siebert, J., Spencer, R., and Ray, A. (2002). Validation of a vision-based, three-dimensional facial imaging system. *Cleft PalateCraniofacial Journal 40*(5), 523–529.

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

Bowman, A., Katina, S., Smith, J., and Brown, D. (2012). Anatomical curve identification.

Bowman, A. W., Katina, S., Smith, J., and Brown, D. (2013). Anatomical curve identification.

Cantzler, H. and Fisher, R. B. (2001). Comparison of hk and sc curvature description methods. pp. 285–291. Third International Conference on 3-D Digital Imaging and Modeling.

Dryden, I. and Mardia, K. (1998). *Statistical Shape Analysis*.

Eilers, P. H. C. and Marx, B. D. (1996). Flexible smoothing with b-splines and penalties. *Statistical Science 11*(2), 89–102.

Einbeck, J., Tutz, G., and Evers, L. (2005). Local principal curves. *Springer Science 15*, 301–313.

Gairdner, D. and Pearson, J. (1988a). Growth and development record. boys: Preterm2 years length/weight/head circumference.

Gairdner, D. and Pearson, J. (1988b). Growth and development record. girls: Preterm2 years length/weight/head circumference.

Goldfeather, J. and Interrante, V. (2004). A novel cubic-order algorithm for apporiximating principal direction vectors. *ACM Transactions on Graphics 23*, 45–63.

Hastie, T. and Stuetzle, W. (1989). Principal curves. *Journal of the American Statistical Assocation 84*, 502–516.

Kau, C., Zhurov, A., Richmond, S., Bibb, R., Sugar, A., Knox, J., and Hartles, F. (2006). The 3-dimensional construction of the average 11-year-old child face: A clinical evaluation and application. *American Association of Oral and Maxillofacial Surgeons 64*, 1086–1092.

Kent, J., Mardia, K., Morris, R., and Aykroyd, R. (2001). Functional models of growth for landmark data. pp. 109–115. Leeds: Leeds University Press.

Koenderink, J. and Doorn, A. (1992). Surface shape and curvature scales. *Butterworth-Heinemann Ltd 10*, 557–564.

Morris, R., Kent, J., Mardia, K., and Aykroyd, R. (2000). A parallel growth model for shape. pp. 171–174. Bristol: BMVA.

Morris, R., Kent, J., Mardia, K., Fidrich, M., Aykroyd, R., and Linney, A. (1999). Analysing growth in faces. pp. 77–91. Las Vegas: CSREA Press.

NHS (2012). Cleft lip and palate. [Online; updated 05-July-2012].

White, J., Ayoub, A., Hosey, M., Bock, M., Bowman, A., Bowman, J., Siebert, J., and Ray, A. (2003). Three-dimensional facial characteristics of caucasian infants without cleft and correlation with body measurements. *Cleft PalateCraniofacial Journal 41*(6), 593–602.

Yamada, T., Mori, Y., Minami, K., Mishima, K., and Tsukamoto, Y. (2001). Three-dimensional analysis of facial morphology in normal japanese chil-

dren as control data for cleft surgery. *Cleft PalateCraniofacial Journal 39*(5), 517–526.