4

Bone Shapes from a Paleopathology Study

4.1 Archaeology and arthritis

Archaeologists have conducted a major excavation at St. Peter's Church, Barton-upon-Humber, in the north of England. They have exhumed the skeletons of about 2000 adults dating mainly from between 1000 and 1500 C.E. A particular way in which the bones have been studied is for *paleopathology*—the use of old remains to give us information about diseases that people suffered from in the past. Many diseases leave traces on the bones, and special attention was given to osteoarthritis of the knee, both because it is and was a common and painful disease, and because the skeletal remains give us easy access to parts of the knee joint not easily seen on X-rays.

The paleopathologists attempted to identify every person in the sample with definite signs of osteoarthritis of the knee, as evidenced by *eburnation*—polished bone surface caused by complete cartilage loss. Initially, 23 people were found with eburnation on at least one femur. For each such person, controls matched approximately by age, sex, and period of burial were found from among those with no evidence of osteoarthritis at any joint. Once the joints with postmortem damage had been eliminated, this left 16 eburnated femora and 52 controls for analysis.

Several aspects of the biomechanics of the knee have been studied in relation to osteoarthritis. These include obesity, injury, and lower limb malalignment, but the shape of the joint itself has not been very much considered. It has been hypothesized that osteoarthritis can affect bone

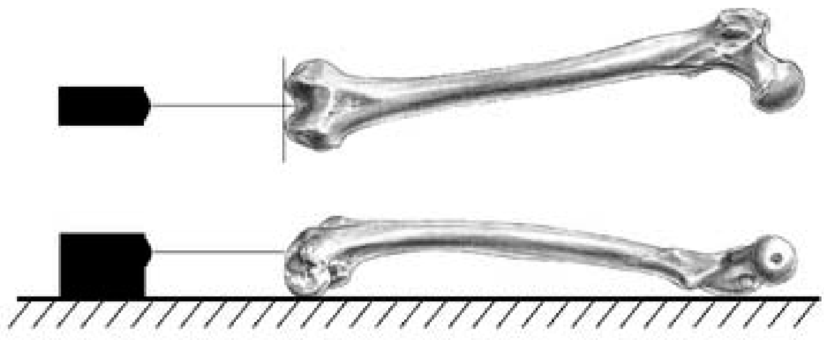

Figure 4.1. Setup showing how the image is captured. A camera captured a digital image of each bone. From Shepstone et al. (1999).

shape, or conversely that certain joint shapes may affect the biomechanics of the joint and hence increase the risk of osteoarthritis. It is against this background that a study of the shapes of the bones was carried out.

4.2 Data capture

As is typical, the investigation had to be carried out rapidly and with a low budget, and so it was not possible to study the three-dimensional structure directly. However, very interesting conclusions can be drawn from simpler two-dimensional images of the joint shape. The first step was to capture the data themselves. Each bone was photographed end-on, as in Figure 4.1, to yield an image as shown in Figure 1.5.

The easiest way of identifying the shape of the joint was to "mark up" each image on the screen by direct reference to the actual bone. The result was a pixel image, with certain pixels specified as being within the outline of the joint. All left femora were reflected to produce "right" images, in order to give every bone a consistent orientation. A typical image is shown in Figure 4.2. The knee end forms an inverted U-shape. The two arms of the inverted U-shape formed by the knee are called *condyles*, and the space between them is the *intercondylar notch*. The smaller indentation at the top of the image is called the *patellar groove*.

For our analysis, we have 68 outlines, of which some are known to correspond to arthritic joints. We regard each outline as a single data object, and consider ways of studying the variability in shape between the bones, and of relating this variability to the presence or absence of arthritis. The first step in studying the shapes is to parameterize the images in an appropriate way. One way of doing this is by defining landmarks; these give a natural way of representing a shape by a fairly low-dimensional array of

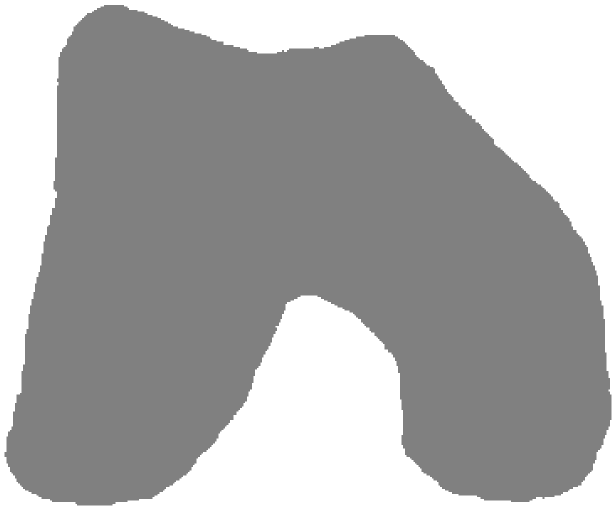

Figure 4.2. Bitmap image after drawing round the outline in Figure 1.5 on the screen and reversing to yield standard orientation.

numbers. In Chapter 8 we return to these data and consider a different approach concentrating on the intercondylar notch alone.

4.3 How are the shapes parameterized?

The principle of using landmarks is to locate a fairly small collection of points from which the shape itself can be reasonably reconstructed. The process used for the bone shapes is best described by reference to Figure 4.3. Initially, the landmarks numbered 1, 2, 5, 7, 9, and 12 were located 'by hand' (in fact by mouse) on the image. These correspond to lowest and highest points on the relevant part of the outline, but because of the strange shapes of some of the specimens, are easier located manually than algorithmically. Then landmarks 3, 6, 8, and 11 were defined as the extreme points within the image of the perpendicular bisector of the lines 2–5, 5–7, 7–9 and 9–12 respectively. This process was repeated on the lines 3–5 and 9–11 to give landmarks 4 and 10. For the remainder of the analysis, we discarded the bone pixel images and worked with the landmarks.

Any bone's shape can be reasonably well approximated by putting a smooth curve through the coordinates of the 12 landmarks. Although the calculations we carry out are in terms of the 24 coordinates of the landmarks, conceptually we are considering the shapes as the data of interest,

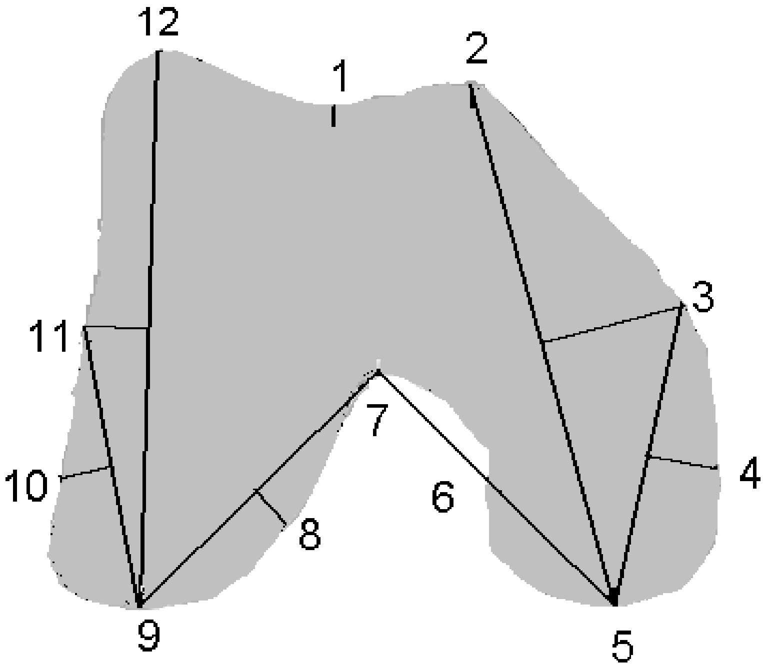

Figure 4.3. Demonstration of the process of identifying and constructing landmarks. Landmarks 1, 2, 5, 7, 9, and 12 are located manually, and the others are then found automatically, as the extreme points within the image of the perpendicular bisector of the lines shown. From Shepstone et al. (1999).

and the results in terms of the curves themselves. To each set of landmark positions there corresponds a periodic curve, and the coordinates of the landmarks are the way that the curves are represented internally to our calculations.

To be precise, the interpolation is carried out by fitting periodic cubic spline interpolants to the landmark x and y values separately, to give functions $x(t)$ and $y(t)$ for t in $[0, 1]$. A cubic spline is a curve made of pieces of cubic polynomials, joined together smoothly at the data points, and the fitting was done using the S-PLUS routine `spline`. The landmark positions gave the values of x and y at the points $i/12$ for $i = 1, 2, \ldots, 12$. As t varies, the point $(x(t), y(t))$ then traces out the curve. The same technique is used whenever we wish to recover a curve from its landmark positions. In mathematical terms, continuing the ideas discussed in Section 2.5, we have implicitly constructed a basis for the representation of these shapes.

4.4 A functional principal components analysis

4.4.1 Procrustes rotation and PCA calculation

Because the size and orientation of the bones is of no particular interest, we eliminate size and orientation variability by a process known as *Procrustes transformation.* In Greek mythology, Procrustes was a robber who captured passing travelers and made them fit his bed, either by stretching their limbs or by chopping them off. Fortunately the analysis of data is less traumatic, but the idea is still to adjust the data so they fit together as closely as possible. First each configuration is centered at its mean, in order to eliminate any translation effects. Then the configurations are all rotated and scaled to minimize the sum of squares between the configurations. For software details, see the Web page associated with this chapter.

Let $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots, \boldsymbol{\mu}_{12}$ be the mean positions of the 12 landmarks, after transformation. Let μ be the interpolating curve between these positions, constructed in the way described in Section 4.3. Then μ is considered as the mean bone shape.

Each individual shape yields a vector of 24 coordinates, the x and y coordinates of the 12 landmarks. (Because of the Procrustes fitting, there are some dependences between these coordinates, but that does not affect the subsequent work.) We perform a functional principal components analysis of the 68 curves by using standard principal components analysis on the 68 24-vectors of landmark coordinates. Before examining the results, it is worth reviewing the way in which this functional principal components analysis can be interpreted.

4.4.2 Visualizing the components of shape variability

Concentrate first on the leading component. For this component, standard PCA provides a 24-vector of principal component loadings, which can be expressed as twelve 2-vectors $\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_{12}$. As we saw in Section 2.3, a good way of visualizing the relevant variation is to plot curves corresponding to the mean plus and minus a multiple of the effect of variation in this component direction. Indeed, in the shape context it is hardly meaningful to consider the principal component weights aside from their effect on a particular shape such as the mean. In the present example, three standard deviations of the principal component give a suitable multiple; more generally the choice may have to be adjusted subjectively.

Let s be the sample standard deviation of the principal component. We then find two curves, plotted in Figure 4.4. The solid curve is the interpolant to the landmarks $\boldsymbol{\mu}_1 + 3s\mathbf{z}_1, \boldsymbol{\mu}_2 + 3s\mathbf{z}_2, \ldots, \boldsymbol{\mu}_{12} + 3s\mathbf{z}_{12}$. The first principal component of this curve will be $3s$, and it will exemplify the kind of curve that has a positive value of the first principal component. The dashed curve is the interpolant to $\boldsymbol{\mu}_1 - 3s\mathbf{z}_1, \boldsymbol{\mu}_2 - 3s\mathbf{z}_2, \ldots, \boldsymbol{\mu}_{12} - 3s\mathbf{z}_{12}$, and will have

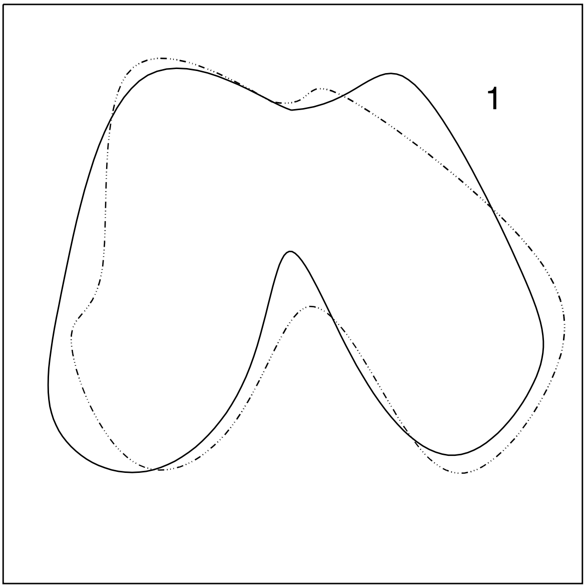

Figure 4.4. The effect of the first principal component of variation. The curves correspond to the mean $\pm$ three standard deviations of the component. The solid curve is the effect of adding the component and the dashed curve of subtracting it. This component explains 21% of the variability in the original data.

a negative value of the first principal component. Furthermore, the two curves indicate the variability of the first principal component within the data, because of the choice of a multiple depending on s. In the present case we do not plot the mean shape itself, because the mean can be inferred by eye from the given curves.

It can be seen from Figure 4.4 that if an outline has a positive score on the first principal component, then we can expect it to have a deeper intercondylar notch, and also a more pronounced bulge in the top right part of the image. The converse characteristics would be associated with a negative value of this component.

Similar plots for each of the principal components 2 to 5 are shown in Figure 4.5. The second component will be of particular importance; a positive score is associated with a narrowing of the right-hand condyle (in our diagram) and with a deepening and widening of the intercondylar notch.

How do arthritic bones differ from controls? For each component, a t-test was carried out to compare the eburnated and noneburnated bones. There was no significant difference on components 1, 3, 4, and 5, but the difference on component 2 was highly significant ($t = -3.01$, $p = 0.0037$). On this component, the mean for the eburnated bones was -10.9 and for

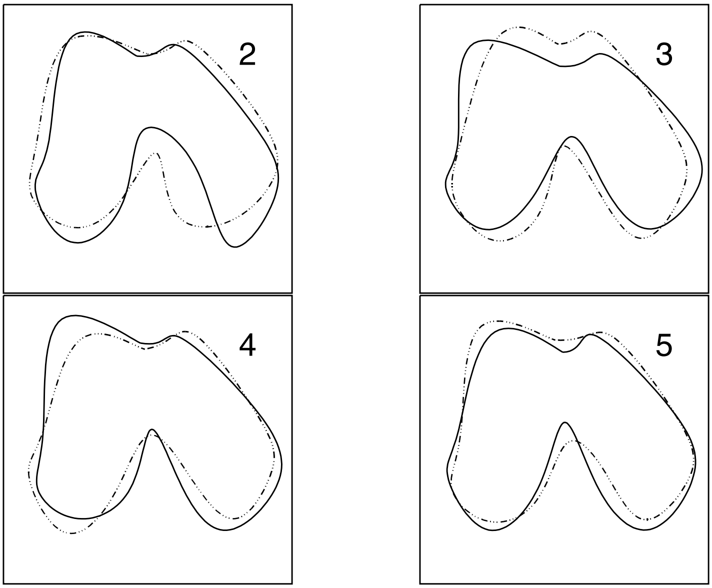

Figure 4.5. The effects of the second to fifth principal components of variation. These explain 18%, 12%, 9%, and 8% of the original variability, respectively. Only on the second component is there a significant difference ($p = 0.0037$) between the eburnated and noneburnated bones. On this component the mean score for the eburnated bones was significantly higher than for the controls.

the controls it was 3.4. This indicates that, on the average, the eburnated bones will tend to have the properties associated with a positive score on component 2.

4.5 Varimax rotation of the principal components

It is well known in classical multivariate analysis that an appropriate rotation of the principal components can, on occasion, give components of variability more informative than the original components themselves. A rotation method constructs new components based on the first k principal components, for some relatively small k. The idea is that k is chosen to include all the components that convey meaningful information, but not those that are just “noise”. In the present example, we concentrate on the first five components and set $k = 5$.

The varimax method is often a useful approach. The method chooses components to maximize the variability of the *squared* principal component weights. The resulting modes of variability tend to be concentrated on part of the range of the function in question, so in the present context

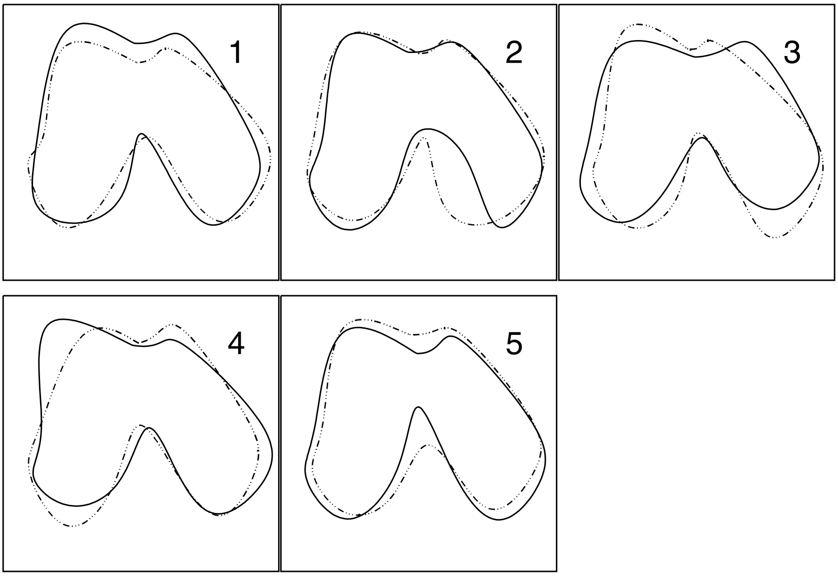

Figure 4.6. The effects of the first five varimax-rotated components for the bone shape data. The percentages of variability explained are, respectively, 14%, 15%, 15%, 11%, and 13%. The arthritic bones had significantly higher scores than the controls on component 2 and significantly lower on component 3.

they express departures from the mean curve over part of the outline shape rather than the whole of it. They are still orthogonal, but the values of the components for the data will no longer necessarily be uncorrelated. Furthermore, the variances of the varimax components will be less spread out than those of ordinary components, and need no longer decrease monotonically. The varimax algorithm is discussed further in Section 4.8.

The modes of variation corresponding to the varimax-rotated components are shown in Figure 4.6. Compared to the original principal components in Figures 4.4 and 4.5, some of the varimax components are more definitely interpretable in terms of the bone shape. Varimax component 2 completely corresponds to a thinner right condyle, in the orientation shown in the figure. Component 5 is concentrated almost entirely on a much narrower join between the condyles. Component 3 is associated with a broader intercondylar notch, but more particularly with a much more symmetric patellar groove than the mean.

The percentages of variances explained by the components are roughly the same for each of the components displayed. As with the raw principal components, the component scores for the two classes of bones were compared. On components 2 and 3 the difference is significant, but not as strongly as previously ($p < 0.025$ in both cases). On component 2 the eburnated bones tend to have negative scores, whereas their scores on component 3 tend to be larger than average. This suggests that the eburnated bones tend to have a thicker right condyle, and a flatter and more symmetric patellar groove.

Is varimax rotation worthwhile? It yields components that have much more direct meaning for the bone shapes themselves. In terms of finding ways in which the two groups of bones differ, it highlights two components rather than concentrating attention on a single component. However, the individual interpretation of each of these two components, especially varimax component 2, is much more physically intuitive than the composite effect represented by original component 2 in Figure 4.4.

4.6 Bone shapes and arthritis: Clinical relationship?

The relationship between the shape of the femur and the incidence of osteoarthritis of the knee has not been studied widely, and so any clinical conclusions have to be tentative. It is possible to analyze the data further, for example, by breaking down the eburnated group according to the position of the eburnation. There is then some suggestion that the location of the eburnation is associated with the third varimax-rotated component score, corresponding to the variation in shape of the patellar groove. On the other hand, the change in shape of the condyles associated with the second varimax component seems only to be associated with presence or absence of eburnation. However, the numbers of bones in each subgroup are not sufficient to draw firm conclusions.

What is the possible link between arthritis and the shape of the bones? On the basis of these data alone, it is not possible to discover to what extent shape variation in the condyle is a cause or an effect of osteoarthritis. Differences in intercondylar notch shape could conceivably affect the functioning of the ligaments in the joint, or increase the likelihood of damage, and lead to an increased risk of knee osteoarthritis. Conversely, arthritis causes a change in biomechanics, which could possibly lead to bone remodeling. An increase in the width of the condyle would help to stabilize an unstable joint or dissipate increased pressure. The data support the concept of a feedback mechanism within which this kind of reshaping of joints is an attempt to slow, or counter, the effects of osteoarthritis.

The association of eburnation with the shape of the patellar groove is more of a puzzle. Postmortem studies have shown a naturally occurring wide variation in patellar groove shape. This could be a potential risk factor, with a wide and shallow groove leading to biomechanical differences that can cause osteoarthritis. However, the potential mechanisms are not yet well understood.

4.7 What have we seen?

Functional data do not have to be a simple function of one variable, but can take many other forms. For the analysis of the bone shape data, the functions of interest were shapes as described by cyclic curves in two dimensions. An interesting topic for future research would be the consideration of the full three-dimensional joint shape; leaving aside statistical considerations, in the present context this would have been impossible because appropriate data-collection equipment was not available.

Landmarks can provide a very good way of representing functional data. We think about our data as functions, but we have to represent them in a finite-dimensional way in order to carry out calculations, and landmarks are one way of getting a finite-dimensional representation. The landmarks may or may not be of direct interest in themselves—in this chapter they were only the means to the end of considering the function as a whole.

Principal components analysis gave us the way of identifying important modes of variability in the data. In some data sets we would study the values of the principal components on individuals, but in this case it was of particular interest to compare two groups, the eburnated and the noneburnated bones. Once the principal component scores had been found, standard statistical techniques could be used to compare the groups.

The varimax procedure improved the interpretability of the components to some extent, and also was useful in subsequent analysis taking into account the position of eburnation. Varimax and other rotation methods are not a panacea, but will often provide a helpful contribution to the analysis of the data.

4.8 Notes and bibliography

Much of the material of this chapter is based on Shepstone, Rogers, Kirwan, and Silverman (1999), although the method of varimax rotation is somewhat different. They give details of data collection and of the arthritis background, with many references to the relevant literature. They also give further discussion of the conclusions drawn in Section 4.6 above. The data collection from the original bones was carried out as part of Lee Shepstone's Ph.D. research (Shepstone, 1998), under the supervision of the other three authors of the paper.

The use of landmarks to characterize curves is discussed in Ramsay and Silverman (1997, Chapter 5). Dryden and Mardia (1998) give a full discussion of landmark-based methods of the analysis of shape data, together with many references to the literature on statistics of shape. For more material and references on functional PCA, see Ramsay and Silverman (1997, Chapter 6). Their Section 6.3.3 gives some discussion of the varimax-rotation procedure. Fuller details of varimax rotation, and also of Procrustes fitting, are given in standard multivariate text books. See, for example, Harman (1976), or Mardia, Kent, and Bibby (1979).

There is a subtle point to be taken into account in the case we have discussed. Each landmark is a 2-vector, and so the principal component weights are themselves 2-vectors. We therefore base the varimax criterion on the variability of the squared lengths of the 2-vectors of principal component weights, rather than directly on the individual weights. Suppose that the loadings of the first five principal component weights are given by two 12×5 matrices $\mathbf{A}^X$ and $\mathbf{A}^Y$, respectively containing the loadings on the x- and y-coordinates of the 12 landmarks. We aim to find a 5×5 rotation matrix $\mathbf{T}$, yielding rotated loadings matrices $\mathbf{B}^X = \mathbf{A}^X\mathbf{T}$ and $\mathbf{B}^Y = \mathbf{A}^Y\mathbf{T}$, to maximize the variance of the quantity

$$\sum_{i=1}^{12}\sum_{k=1}^{5} ||b_{ik}||^2,$$

where b_{ik} is the 2-vector $(\mathbf{B}^X_{ik}, \mathbf{B}^Y_{ik})$. See the Web page for this chapter for further details.