Chapter 6

Diaphysis segmentation

In Section 2.5.1 on page 25, long-bone diaphyses were identified as the location of interest for detecting fractures. Since both the anatomic structure and function of a long-bone varies depending on the location examined, the image texture also varies along the length of the bone. To simplify the fracture detection task, the goal was to segment the bone into separate regions based on these texture differences, and to use a different fracture detection technique for each texture. Rather than analysing texture to locate the boundaries between the regions, the AO scheme presented in Section 2.5.1 on page 25 and shown in Figure 2.10b on page 18 was utilised. The steps required to segment the diaphysis from an AP x-ray image have been implemented in a new algorithm that is described in this chapter. The AO segmentation method suffered from a number of limitations that are examined in Section 6.4. Consequently a second method of segmenting the diaphysis based on analysing the bone curvature was introduced, and its performance compared to the AO method.

6.1 AO algorithm overview

Figure 6.1 shows how the diaphysis is segmented from the image, and each step of the process is detailed in the following sections. Briefly, the first step in automatically locating the diaphyseal boundary using the AO method is to determine the bone centre-line (shown in Figure 6.1a). Once the centre-line parameters are determined, the extreme articular surface—that is the extreme endpoint of the bone being segmented—

must be located (Figure 6.1b). The maximum epiphyseal width $\omega_{epiphysis}$ of the bone can then be measured (Figure 6.1c) and used to mark the diaphyseal boundary as being that same distance $\omega_{epiphysis}$ from the articular surface (Figure 6.1d). Each of these segmentation steps are discussed in detail in the following sections, and performance results for the algorithm are also presented.

6.2 Bone centre-line estimation

The first segmentation stage used the long-bone approximation parameters (ρ_i, θ_i), determined using the methods described in Chapter 5, to calculate the best estimate of the bone centre-line. The bounding box, and therefore the diaphyseal boundary, could not be arbitrarily oriented, but had to be correctly aligned with the bone axis. Accordingly, it was critical that the bone centre-line (or centre-lines) be correctly determined, since all subsequent AO scheme measurements were made either parallel or perpendicular to this line. In some cases, the location of the bone centre-line was relatively easy to determine, while in other cases it was much more complicated. In general, centre-line detection was much more difficult for displaced fractures than undisplaced fractures, and also for distal bones than proximal bones. Undisplaced fractures were those in which the centre-lines of the shaft segments on either side of the fracture remained colinear, while the centre-lines of the shaft segments in displaced fractures did not remain colinear. This meant that centre-line detection was hardest in displaced distal fractures.

The estimation of the appropriate bone centre-line was split into two distinct steps:

1. Estimation of the centre-line(s) of all the relevant bones in the image, using the information obtained in Chapter 5.

2. Selection of the most appropriate centre-line(s) from the complete set of centre-lines found in step one.

The methods by which both of these steps were performed are discussed in the following sections.

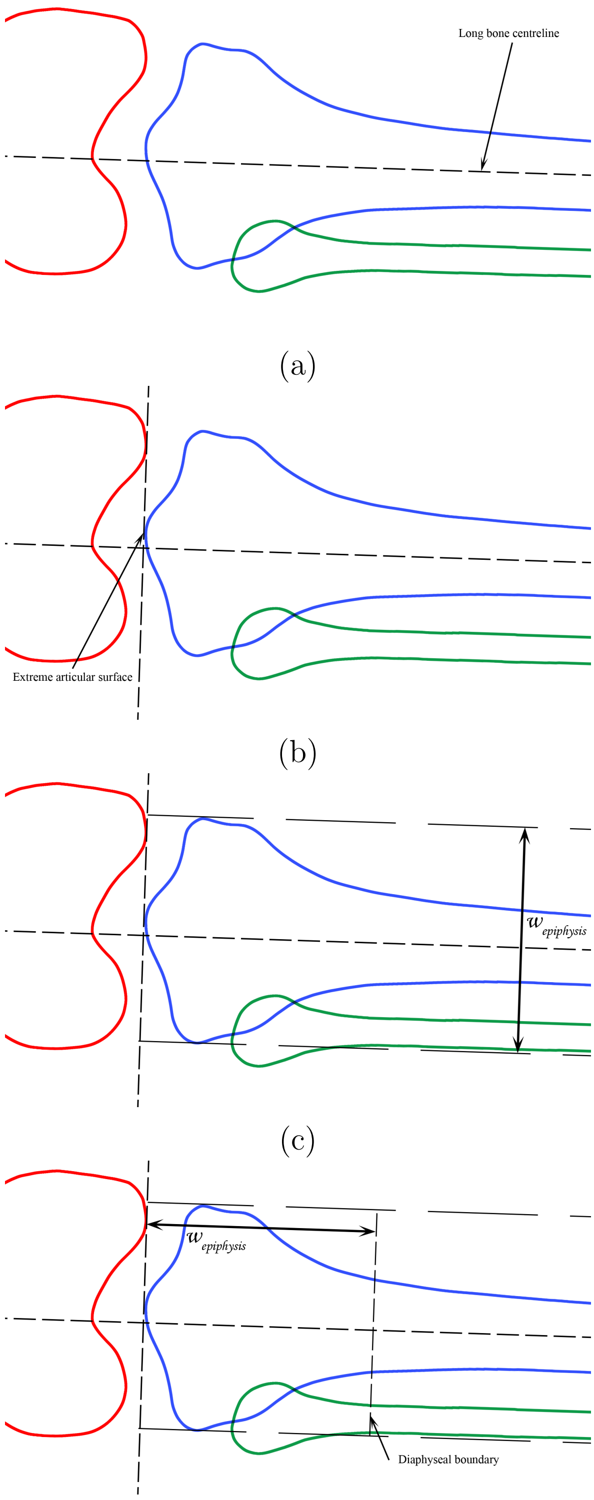

(a)

(b)

(c)

(d)

Figure 6.1: An example of the steps and measurements required for diaphysis segmentation. (a) First the bone centre-line is determined, and (b) the extreme articular surface located. (c) Then the epiphyseal width is measured, and then (d) this is used to mark the diaphyseal boundary.

6.2.1 Locating all the bone centre-lines

To simplify the task of centre-line detection, the line segments located in Section 5.7.3 were utilised. Assuming that these segments were correctly detected, the bone centre-lines should be found by pairing the e identified lines, because each centre-line should be located somewhere between the bone edges approximated by those lines. The next problem was determining the criteria by which the peaks should be paired. Simple solutions such as pairing the lines based only on θ or ρ did not always perform well, especially in images of displaced fractures where the bones were offset. This also meant that pairing the detected lines by purely analysing the peaks in the Hough parameter space also produced poor results. This was because there were many possible accumulator peak patterns to identify, and also because again the extent of the lines was unknown (the same problem that was identified in Section 5.7). Therefore the centre-line detection was best performed in the image space where the line endpoints were known, and needed to utilise a robust pairing method to correctly pair the identified peaks. The following subsections discuss some different methods of detecting the centre-line or centre-lines of a long-bone.

Averaging the parameters

X-ray images of unfractured proximal long-bones (i. e. the humerus and femur—bones 1 and 3 in Figure 2.10b on page 18) theoretically contained only two closely aligned bone edges, and therefore only two strong peaks in the Hough accumulator. This statement was only true provided the bones were either not fractured, or contained only undisplaced fractures. X-rays of unfractured proximal long-bones were the simplest of all images for calculating the bone centre-line.

The bone centre-line was defined as the line that lay halfway between the two proximal long-bone edges. Accordingly, once the two Hough accumulator peaks were identified using the method described in Section 5.4.1, the centre-line parameters were then calculated by determining the averages of the two θ and two ρ values, thus producing an average angle and average distance from the centre of the image:

$$\theta_c = \frac{\theta_1 + \theta_2}{2} \qquad \rho_c = \frac{\rho_1 + \rho_2}{2} \tag{6.1}$$

Alternatively, the endpoints of the centre-line could be determined by calculating the midpoints of the lines joining the line endpoints:

$$(x,y)_{start} = \left(\frac{x_1 + x_2}{2}, \frac{y_1 + y_2}{2}\right) \qquad (x,y)_{end} = \left(\frac{x_3 + x_4}{2}, \frac{y_3 + y_4}{2}\right) \tag{6.2}$$

where (x_1, y_1) and (x_2, y_2) were the endpoints of the first line, and (x_3, y_3) and (x_4, y_4) were the endpoints of the second line. Once the parameters θ_c and ρ_c, or $(x,y)_{start}$ and $(x,y)_{end}$ were determined, the centre-line could be displayed on the x-ray image by back-projection. An example of bone centre-line detection in a femur x-ray using equation 6.2 is shown in Figure 6.2. A visual examination shows that the dashed green line in Figure 6.2c is indeed a good representation of the femur centre-line.

Although this method produced a good representation of the femur centre-line, its applicability was limited to simple cases. Complicated peak pairing was neither performed nor required because only $e = 2$ peaks were present and therefore had to be paired together. Obviously if a larger number of peaks were present in the accumulator and the parameters of all of those peaks were averaged, then the result was unlikely to be a good representation of the required centre-line. As a result, this method was only suitable for complete centre-line detection in undisplaced proximal long-bones. Once the peaks were appropriately paired in the following methods, Equations 6.1 and 6.2 were then still appropriate for determining the centre-lines of each of the paired lines.

Ranking the ρ or θ parameters

The parameter averaging method did not need to determine which lines were most appropriate to pair because only $e = 2$ peaks were present in the accumulator. In those cases where there were more than $e = 2$ bone edges—such as the distal long-bones or images that contained displaced fractures that resulted in mal-aligned bone segments—the most appropriate method of pairing the peaks needed to be determined. Possible

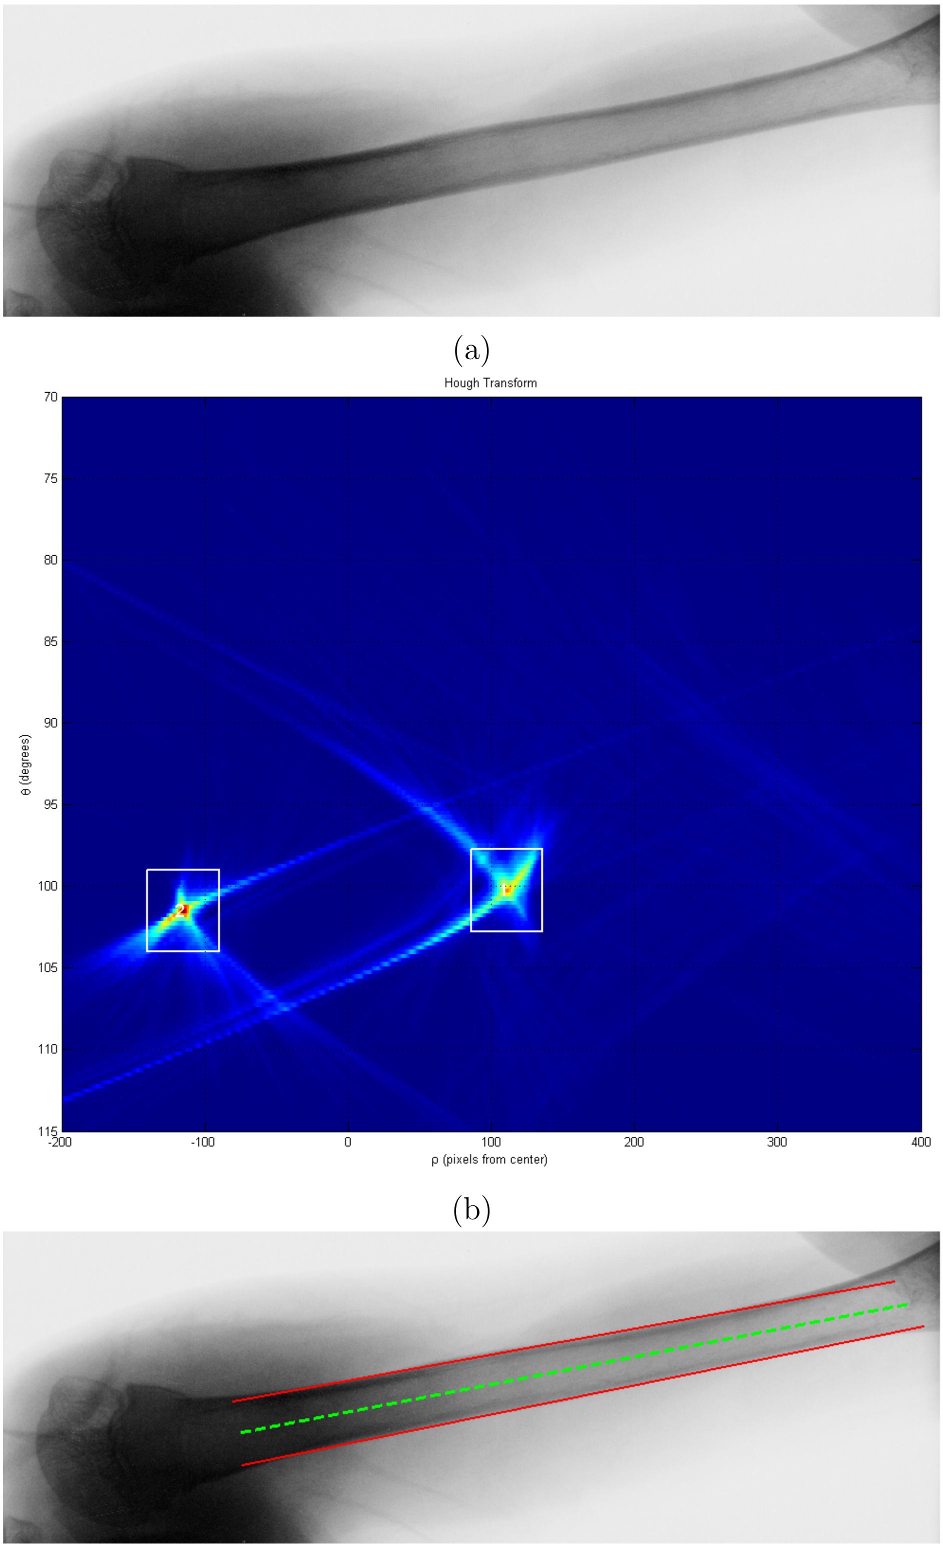

(a)

(b)

(c)

Figure 6.2: (a) An example of an unfractured humerus (subject details unknown), (b) the peaks detected in the Hough Transform, and (c) the corresponding approximated bone edges (red), and the detected centre-line (green).

solutions included ranking the peaks based on either their ρ or θ values.

X-ray images of unfractured distal long-bones (i. e. the radius and ulna, and the tibia and fibula—bones 2 and 4 in Figure 2.10b on page 18) theoretically contained four closely aligned bone edges, and therefore four strong peaks in the accumulator. Once again, this was also true for undisplaced distal long-bone fractures. By way of example, Figure 6.3 shows a midshaft tibia fracture, and the resulting Hough accumulator with the peaks correctly detected. The four peaks were located at approximately the same θ, and were only differentiated by their ρ values. To determine which lines should be paired, the peaks could be ranked purely on their ρ values. Any small variations in θ were disregarded because the lines were assumed to be almost parallel. After being ranked, the top two lines were chosen as one pair while the bottom two were chosen as another pair, and Equations 6.1 and 6.2 were used to calculate the centre-line parameters and centre-line endpoint coordinates. The result produced by this method is shown in Figure 6.3c, and is a good representation of the required centre-lines.

A different method was required for pairing peaks in an image of either the proximal or distal long-bones, when those bones contained a displaced fracture. That is, a fracture in which the bone fragments were displaced or misaligned. Again by way of example, Figure 6.4 shows an example of a displaced midshaft femur fracture. The resulting Hough Transform contained the two peaks that are expected for a proximal long-bone, plus an additional pair of peaks that were located at a different θ. A visual comparison of Figures 6.3b and 6.4b shows that although both contain four peaks in the Hough accumulator, the pattern of the peaks is very different. In this situation the ρ values could not be used to match the peaks because the degree of displacement of the segments caused these values to differ. Instead, to detect the centre-lines, the θ values had to be correctly paired before Equation 6.1 was applied to each pair. For displaced fractures, the peaks could be paired by assuming that the bone edges of a segment should be approximately parallel, regardless of whether the segment was displaced. To determine which edges were most parallel, the peaks are ranked according to their θ value, since lines that were parallel would have similar θ values regardless of how far they were from the centre of the image (that is, regardless of their ρ value). After the

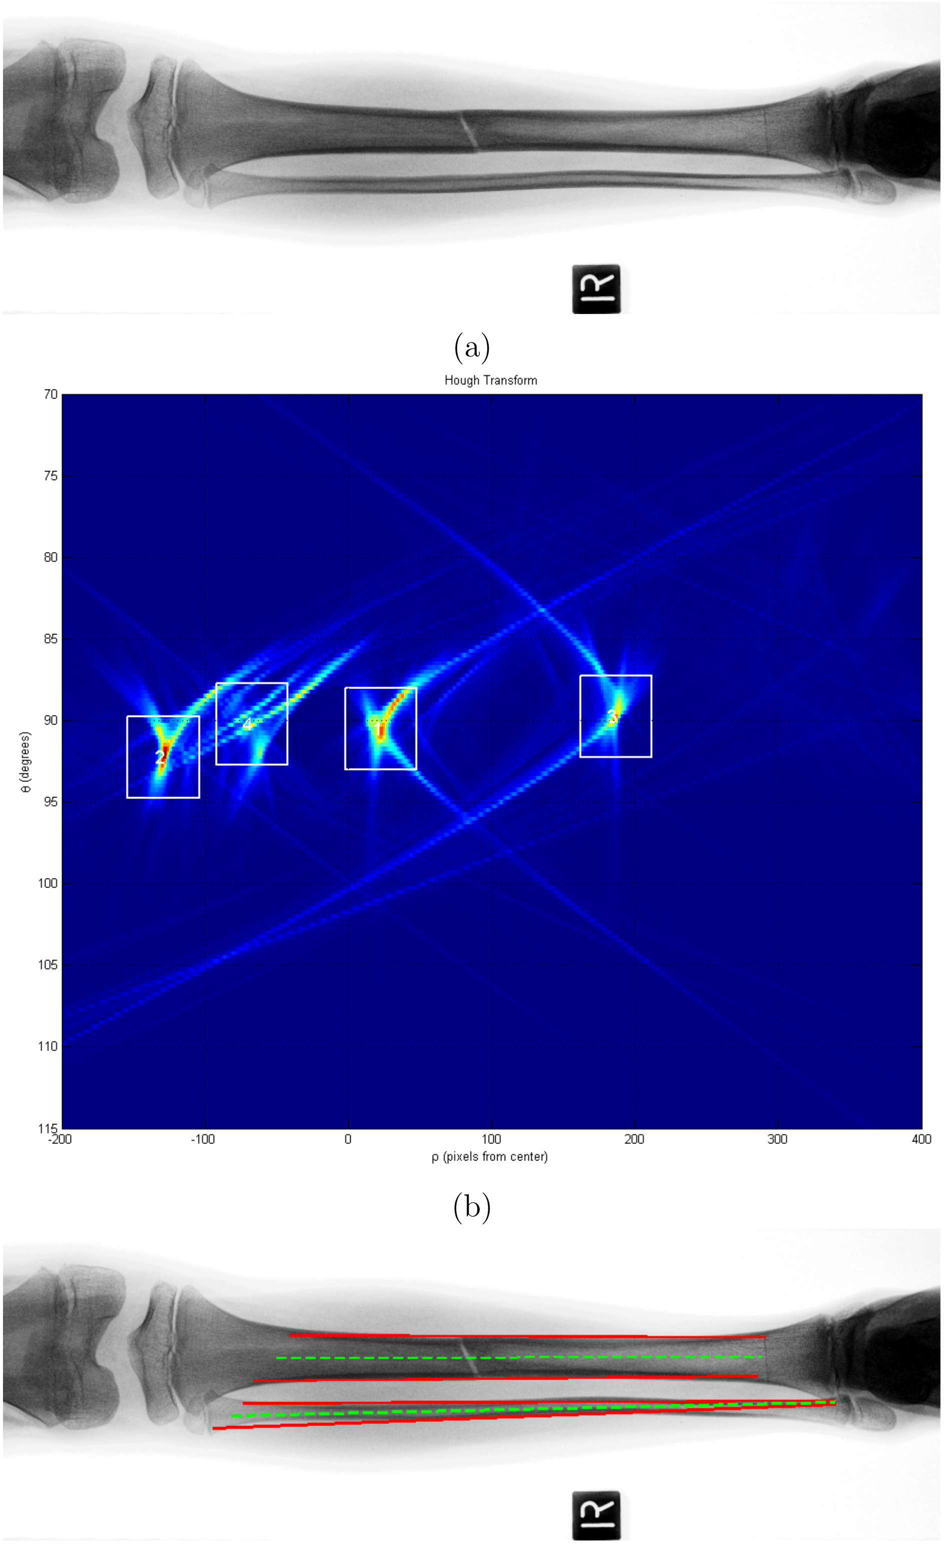

(a)

(b)

(c)

Figure 6.3: (a) An example of an undisplaced midshaft fracture of the tibia in a seven year old male, (b) the peaks detected in the Hough Transform, and (c) the corresponding approximated bone edges (red) and the detected centre-lines (green).

peaks were ranked they were paired in that order, and Equation 6.1 or 6.2 was used to calculate the correct centre-lines.

For example, Figure 6.4b contained peaks at $(\rho,\theta) = (-10, 77.5°)$, $(101, 78.5°)$, $(-8, 64°)$ and $(56, 71.5°)$ so ranking (in this case, in decreasing order of θ) and then pairing produced the two pairs $(\rho,\theta)_A = [(-8, 64°), (56, 71.5°)]$ and $(\rho,\theta)_B = [(-10, 77.5°), (101, 78.5°)]$. The resulting centre-lines $(\rho_c,\theta_c)_A = (39.5, 66°)$ and $(\rho_c,\theta_c)_B = (45.5, 78°)$ were then calculated using equation 6.1, and were marked on Figure 6.4c. Again, a visual examination reveals that the dashed lines are a good representation of the centre-lines of both fragments of the femur.

While ranking the peaks based on either the ρ or θ parameters worked on the two development images shown, there were three reasons why these methods were not a good solution to the problem of detecting all the bone centre-lines. Firstly, while one solution matched the lines based on the ρ values, and the other on the θ values, determining which solution to use was not always a simple choice. While in the examples shown here it was easy to differentiate between the two cases, in many of the more complicated images it was harder to decide which method to use. Secondly, this was compounded by the fact that in some cases, such as displaced distal fractures, it was more appropriate to use a combination of the two methods. Unfortunately when the two methods were combined, the complexity increased and the accuracy of the results decreased significantly. This was especially noticeable when there were, for example, $e = 8$ line segments that had to be correctly paired. Finally, in some cases an odd number of lines were detected during the long-bone shaft parameter approximation stage. This was problematic because the algorithm was then uncertain about which lines should be either paired or excluded, creating incorrect pairings. Thus, a better method of pairing the lines needed to be formulated.

Minimising the distance between the line endpoints

At the start of Section 6.2.1, it was suggested that centre-line detection should be performed in the image space, so that the line endpoint information could be utilised. In the methods previously discussed in this section, this was not the case, because

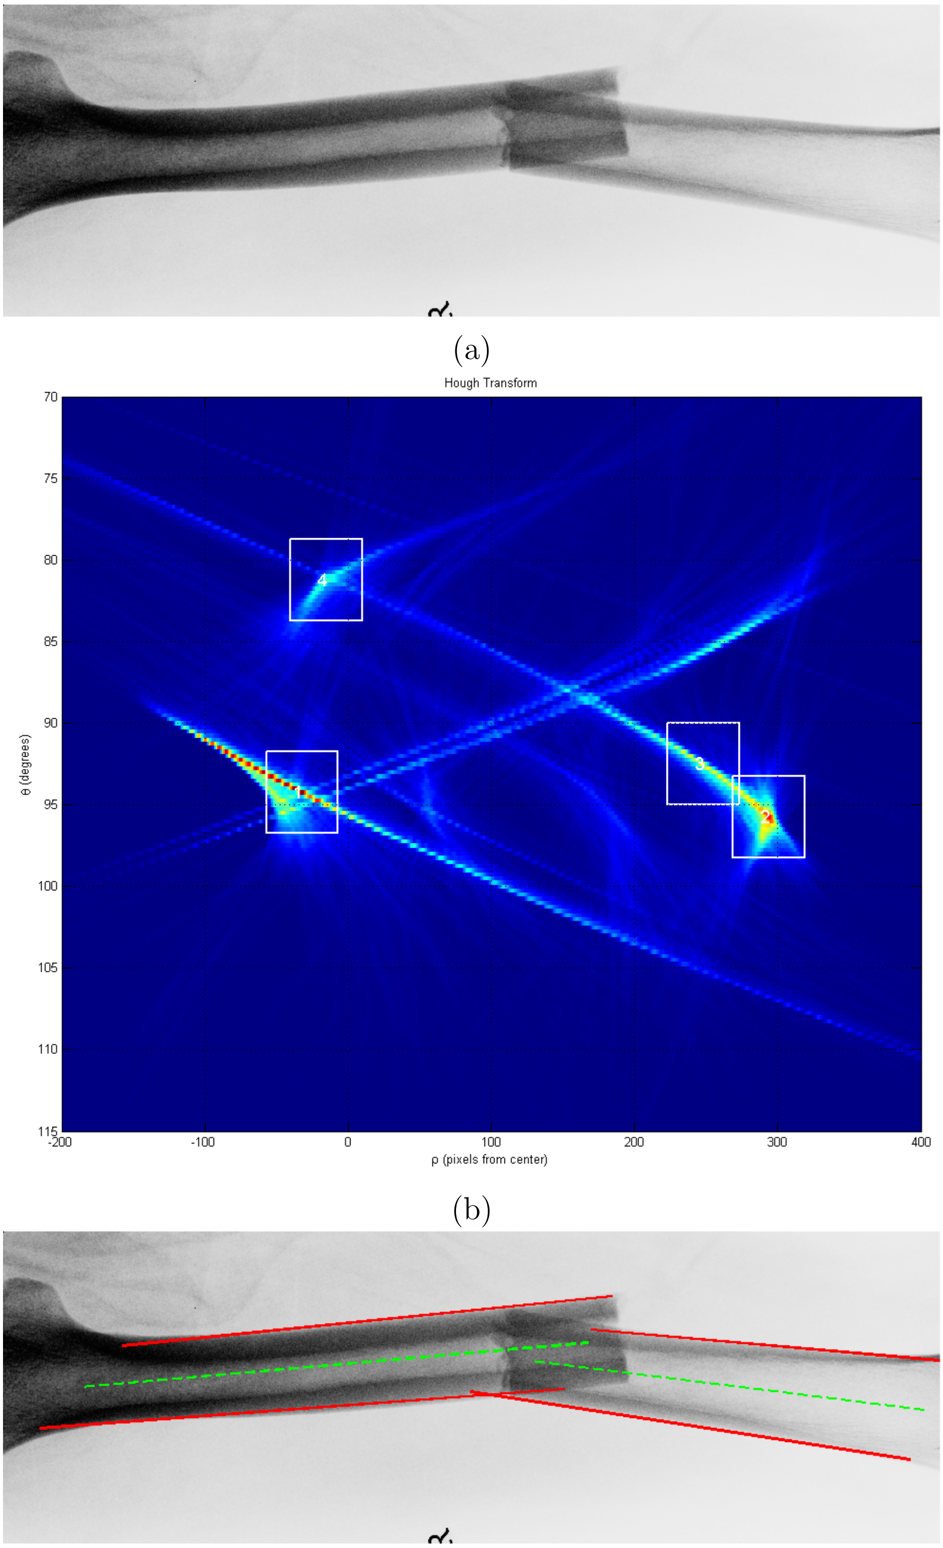

(a)

(b)

(c)

Figure 6.4: (a) An example of a displaced midshaft fracture of the femur in an 18 year old male, (b) the peaks detected in the Hough Transform, and (c) the corresponding approximated bone edges (red) and the correctly detected centre-lines (green).

the analysis of ρ and θ was effectively being performed in the Hough space. This contributed to the poor results produced by both the ρ and θ ranking methods. In contrast, the following analysis was performed in the image space to achieve better results.

Visual inspection of the solid red lines shown in Figures 6.3c and 6.4c, showed that the lines that were to be paired together should be located relatively close to each other, because the bone edges should in most cases be approximately parallel. This also meant that when correctly paired, the line endpoints should also be relatively close together. If an alternative pairing combination was chosen for either of these images, the distances between the line endpoints would be significantly larger. As a result, it was proposed that the criterion for determining the correct pairings would be the minimum total distance between the line endpoints.

To determine the correct pairings, it was necessary to examine every possible combination of line pairs, and record the total distance between the line endpoints. Depending on the number of lines present in the image, there could be a very large number of possible combinations that must be checked. For example, if there were four lines in the image, then there were $4! = 24$ possible ways in which those numbers could be arranged, similarly for eight lines there were $8! = 40320$ ways in which they could be arranged. That is, for N lines, there were $N!$ arrangements. Fortunately, many of those combinations were redundant when selecting two lines at a time. For example, the pairs $[1,2]$ and $[3,4]$ are the same as the combination $[2,1]$ and $[3,4]$, since the order within the pair was not important because both of these combinations produced the same centre-line. In addition, the order of the chosen pairs was also not important since the pairs $[2,1]$, $[3,4]$ were the same as the pairs $[3,4]$, $[2,1]$. To disregard the order within the pairs it was necessary to divide by $2^{\frac{N}{2}}$, and to disregard the order of the pairs themselves it was necessary to divide by $\left(\frac{N}{2}\right)!$ so that the number of possible pairings was:

$$\frac{N!}{2^{\frac{N}{2}}\left(\frac{N}{2}\right)!} \tag{6.3}$$

As a result, for four lines there were a total of only 3 possible ways in which those lines could be paired together, for six lines a total of only 15 possible ways, and for

1 2	3
1 3	2
2 3	1

(a)

1 2	3 4
1 3	2 4
1 4	2 3

(b)

2 3	4 5	1
2 4	3 5	1
2 5	3 5	1
1 3	4 5	2
1 4	3 5	2
1 5	3 4	2
1 2	4 5	3
1 4	2 5	3
1 5	2 4	3
1 2	3 5	4
1 3	2 5	4
1 5	2 3	4
1 2	3 4	5
1 3	2 4	5
1 4	2 3	5

(c)

1 2	3 4	5 6
1 2	3 5	4 6
1 2	3 6	4 5
1 3	2 4	5 6
1 3	2 5	4 6
1 3	2 6	4 5
1 4	2 3	5 6
1 4	2 5	3 6
1 4	2 6	3 5
1 5	2 3	4 6
1 5	2 4	3 6
1 5	2 6	3 4
1 6	2 3	4 5
1 6	2 4	3 5
1 6	2 5	3 4

(d)

Table 6.1: All the possible combinations in which the lines detected in Section 5.7.3 can be paired to locate the bone centre-line or centre-lines. These combinations are shown for (a) three lines, (b) four lines, (c) five lines, and (d) six lines.

eight lines a total of 105 possible ways. When an odd number of lines were present in the image, every combination of an even number of pairs plus one line by itself was examined. Table 6.1 shows the possible combinations in which the lines could be paired for $e = \{3, 4, 5, 6\}$. Once all the combinations that must be checked were known, the analysis could proceed.

For each of the combinations (i. e. each row in Table 6.1), the two Euclidean distances between the line endpoints of each pair (i. e. each column in Table 6.1) were calculated and the sum of these distances was recorded. The combinations were ranked according to the total distance, and the combination with the lowest sum was chosen as the correct pairing. Unfortunately, the single criterion of minimum total distance between the endpoints was insufficient for ensuring that the lines were correctly paired. For a distal long-bone x-ray, it was possible that the lateral edge of the main bone (i. e. the tibia) and the medial edge of the smaller bone (i. e. the fibula) were paired, because they were so close together. This was dependent on the x-ray view, because at some angles these two bones appeared much closer than at others. To ensure that the correct pairs were matched, a second criterion had to be added.

The second criterion was that the absolute difference between the ρ values of each line within the pair was also minimised. This ensured that overlapping pairs such as the medial tibial edge and the lateral fibular edge, were ranked less favourably. Testing on the development images showed that addition was a suitable method of combining the two criteria, with all 28 detected lines in the six images correctly paired. Application of Equation 6.2 again gave the centre-line, producing 14 correct centre-lines in the development set. An example of a displaced distal forearm fracture is shown in Figure 6.5, and demonstrates that minimising the distance between the endpoints allows the centre-lines of each bone fragment to be correctly detected, even when there are $e = 8$ lines present in the image.

6.2.2 Selecting the most appropriate centre-lines

The second stage of bone centre-line estimation involved determining which of the centre-lines detected in the previous stage should be retained, and which, if any, should be discarded. The proximal long-bone images shown in Figures 6.2c and 6.4c were complete and did not require this step, since the centre-lines that they contained were a good representation of the bone centre-line(s), and were suitable for the next stage of the algorithm. However, the distal long-bone images shown in Figures 6.3c and 6.5c were not complete, since they contained too many centre-lines for the next stage of the algorithm. For the distal long-bones, the centre-line was defined as being the axis of the wider bone, so that multiple centre-lines were not identified unless the bone was fractured. The problem was therefore effectively reduced to determining which pair of peaks corresponded to the wider of the two bones, before applying the same averaging process described in Equations 6.1 and 6.2.

Although there was always some variation in the angle θ of the lines corresponding to the four bone edges (as the detected lines in Figure 6.3c show), this could be ignored and only the distances ρ were compared, since it was their distance from the centre

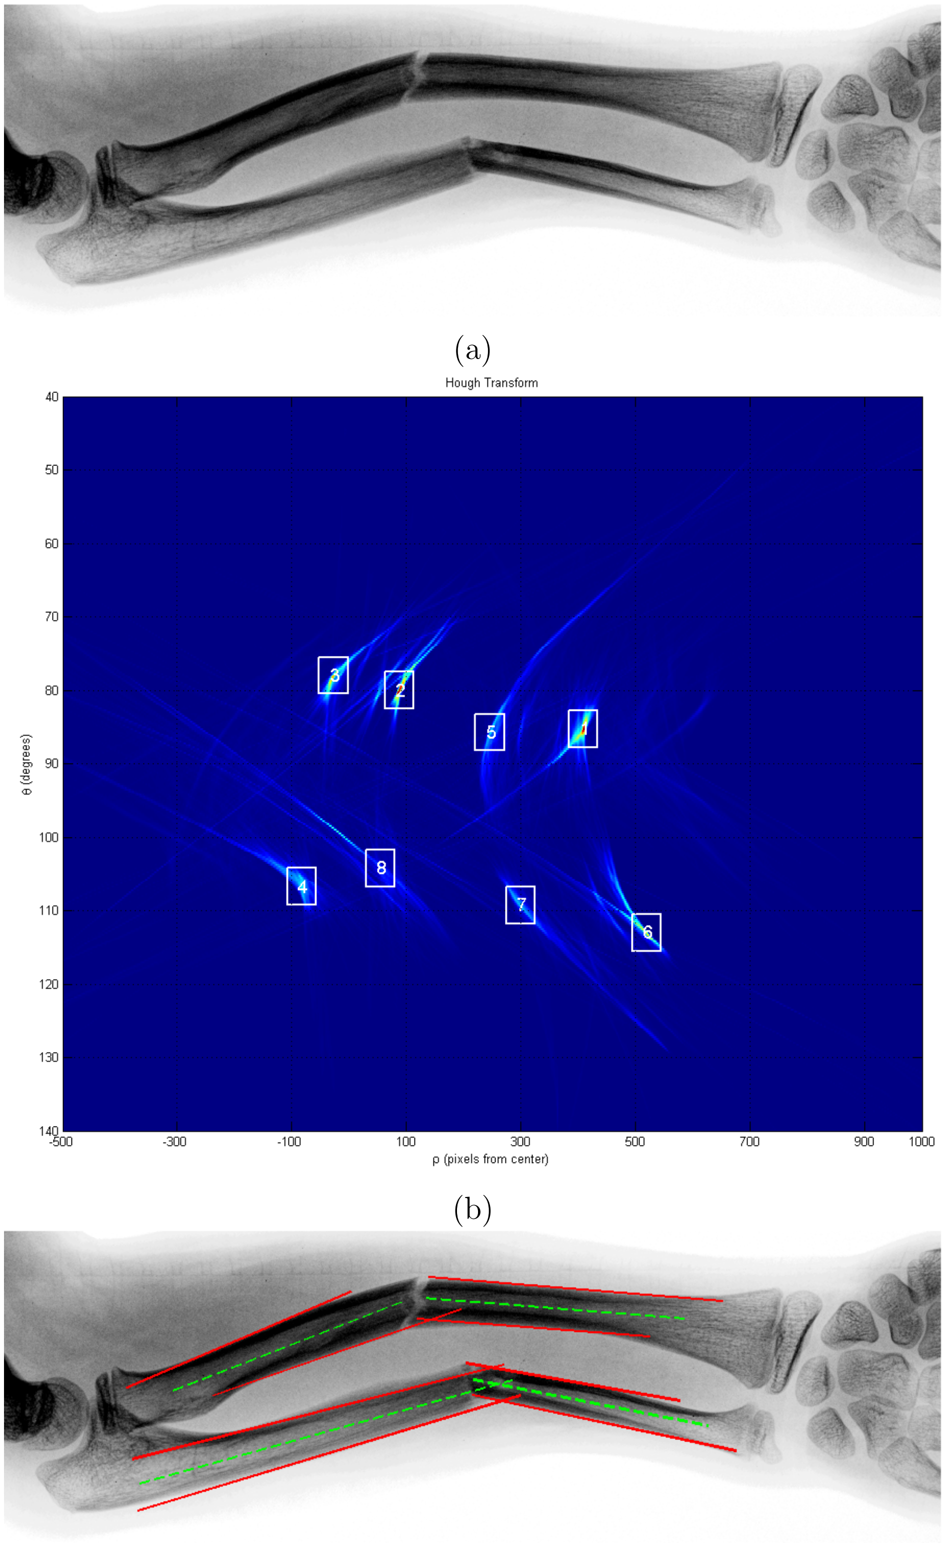

(a)

(b)

(c)

Figure 6.5: (a) An example of a displaced midshaft fracture of the radius and ulna in an eight year old female, (b) the peaks detected in the Hough Transform, and (c) the corresponding approximated bone edges (red) and the detected centre-lines (green).

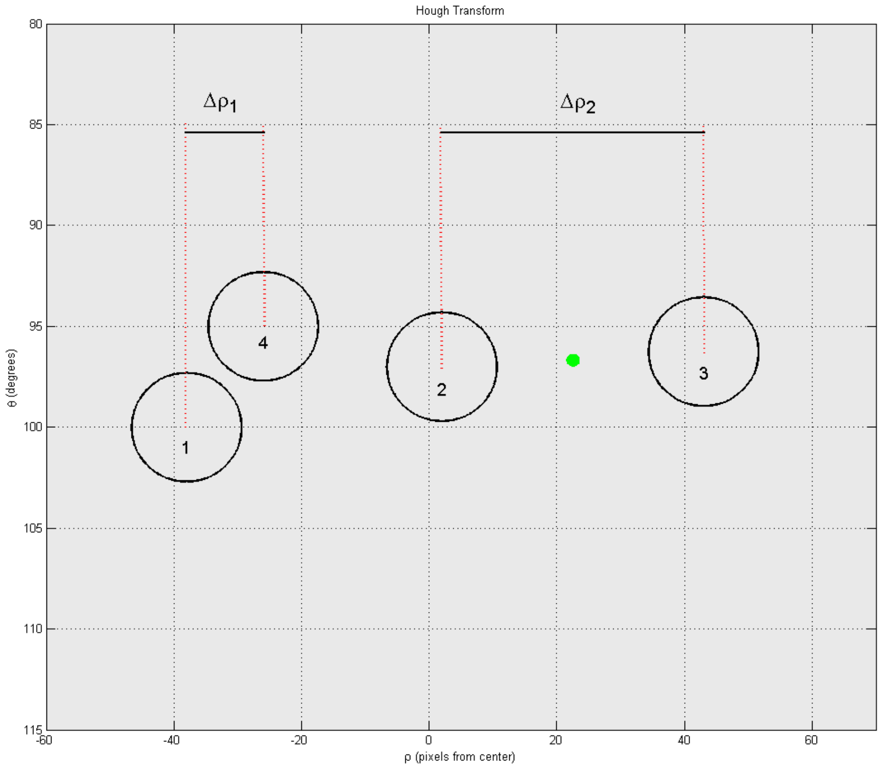

Figure 6.6: The inter-peak distances $\triangle\rho$ in the Hough transform were used to determine which bone was the wider of the pair. Peaks 2 and 3 corresponded to the tibia (the wider bone), while 1 and 4 corresponded to the fibula. Notice that the angular difference of $\triangle\theta \approx 5^\circ$ between the peaks was ignored for this calculation. The green circle represents the chosen centre-line parameters.

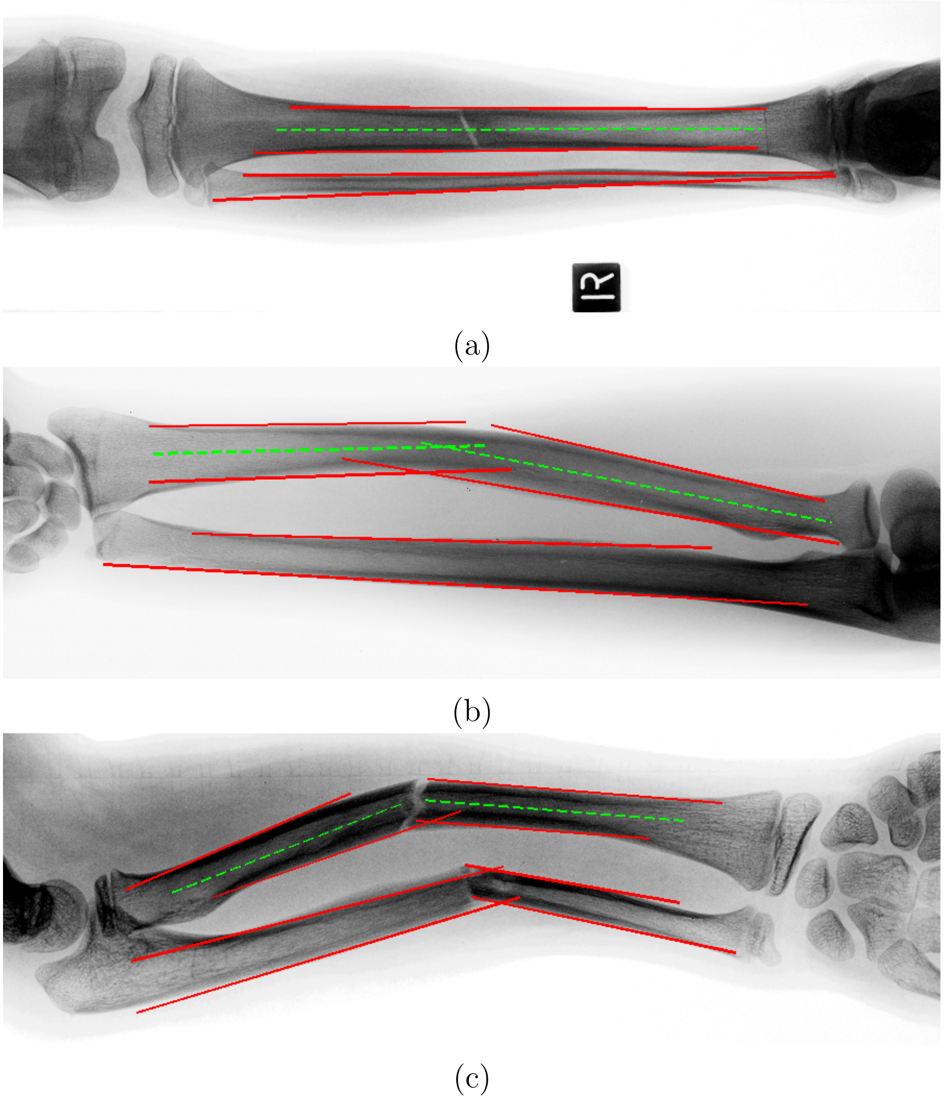

Figure 6.7: Examples of three development images showing the identification of the most appropriate long-bone centre-lines, from the complete set generated in the previous step.

of the image that was important. To correctly pair the peaks, the ρ values were sorted in ascending order and matched with their correct θ values, and the inter-peak distances were calculated using $\triangle\rho_1 = |\rho_2 - \rho_1|$ and $\triangle\rho_2 = |\rho_4 - \rho_3|$. The larger inter-peak distance, $\omega_{shaft} = max\,(\triangle\rho_1, \triangle\rho_2)$, corresponded to the width of the wider of the two bones. As shown in Figure 6.6, the pair of peaks on the right had the largest inter-peak distance. The parameters for the centre-line of that bone were then determined in the same manner as the proximal long-bones (Equation 6.1), while the other centre-line information was discarded. An example from the development set is shown in Figure 6.7a to demonstrate the centre-line that resulted.

In some forearm images—those cases in which six lines were used to approximate the bone due to the curvature of the radius—it was necessary to retain two centre-lines rather than one. In this situation, both portions of the radius required centre-line approximation. The two centre-lines with the most similar ρ value were retained, since they must have corresponded to the same bone, while the other centre-line was discarded. Figure 6.7b shows a forearm image of this type, along with the two detected centre-lines.

Finally, in those images containing displaced distal fractures, a combination of the two methods described here were used. The centre-lines belonging to the same bone were automatically identified by grouping them such that each line in the group had similar ρ values. The widest bone was identified based on the value of ω_{shaft}, and all centre-lines in the group containing that bone were retained while all others were discarded. Figure 6.7c shows the resulting centre-lines for a displaced distal forearm fracture.

6.3 Segmentation using the AO method

After the first segmentation step shown in Figure 6.1a had been performed, and parameters for the centre-line (or centre-lines in the case of the distal long-bones, or fractures containing multiple fragments) had been determined, the second step (Figure 6.1b) could be performed. This involved locating the articular surface of the joint.

6.3.1 Articular surface identification

As described in Section 2.2.1 on page 9 and shown in Figure 2.5 on page 11, the ends of the long-bones are covered by articular cartilage that allows the bones to smoothly slide against each other. The articular surface is three-dimensional, and unlike the simplified joint shown in Figure 2.5a, it is rarely flat. A typical articular surface is much more like that depicted in Figure 2.5b, where the smooth surfaces have a very complicated three-dimensional shape, that varies with anatomic location and also between subjects. Since an x-ray is a projection of a three-dimensional structure onto a two-dimensional image, the complex articular surface appears as a thick edge that traces the end of the long-bone. This section examines how to locate the appropriate point on that thick line.

The point of interest

As described in Section 6.1, the appropriate point on that thick line was termed the extreme articular surface, and it was the furthest most point of the articular surface. If the long-bone was outside the body, and was placed upright on a flat surface so that the identified centre-line was at 90° to that surface, as shown in Figure 6.8, then the point at which the bone sat on the surface was the extreme articular surface. The requirement for locating the extreme articular surface was therefore that the tangent to this surface had to be orthogonal to the bone centre-line, such that:

$$\theta_{tangent} = \theta_{c_i} \pm 90^\circ \tag{6.4}$$

Alternatively, the normal at the extreme point had to have the same θ as the centre-line, but not necessarily the same ρ since the extreme point did not need to occur at the intersection of the centre-line and the articular surface. Equation 6.4 implied that any point on the articular surface that was not orthogonal to the centre-line could not be the most extreme point on the articular surface. Conversely, being orthogonal to the centre-line did not guarantee that point being the extremity, as there may have been a number of points on the articular surface whose tangents were orthogonal to the centre-line. Figure 6.1b and 6.8 show how this extreme point is identified by the tangent orthogonal to the centre-line, and that in both cases this point can be located a small distance away from the centre-line, that is, at a different ρ. Actually identifying the extreme articular surface was equivalent to the problem of determining the location of the furthest tangent to the bone edge that was perpendicular to the centre-line.

Locating the extreme articular surface

Locating the tangent to the bone edge was achieved by examining a projection of the gradient image $|\bigtriangledown I(x,y,t)|$ in the direction $\theta_{tangent}$ orthogonal to the centre-line, and then identifying the appropriate peaks from this projection. Unfortunately without any other processing this method produced unreliable results, because every bone produced a projection containing different features, and identifying the correct peaks was very difficult. However, since the angle of the extreme articular surface $\theta_{tangent}$ was known

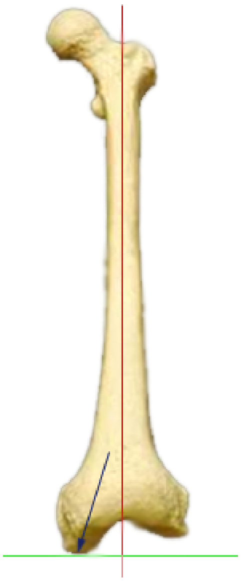

Figure 6.8: If the bone (in this case a femur) is placed upright on a flat surface (green line), such that the centre-line (red line) is vertical, then the extreme articular surface is the point or points (blue arrow) at which the bone touches that surface.

from equation 6.4, the information in $\phi(x,y,t)$ could be used to retain only those points in $|\bigtriangledown I(x,y,t)|$ where the direction of the gradient actually matched the angle of the tangent. That is, the modified Hough Transform described in Section 5.2.3 on page 90 could again be utilised.

Identifying the search region

Figure 6.8 demonstrated that while the extreme articular surface had to lie on the normal to the centre-line, it did not necessarily have to be located on the centre-line itself. However, for all joints other than the neck of femur (bone 3 in Figure 2.10 on page 18), the extreme articular surface was always located close to the centre-line. Hence, the search area could be reduced to a region on either side of the chosen centre-line (ρ_c, θ_c), as shown lightly shaded in Figure 6.9. The size of this region was based on the width of the long-bone shaft ω_{shaft}, as determined in the centre-line calculation using the detected peaks (ρ, θ) from the original Hough Transform. For the proximal long-bones the shaft width was calculated using $\omega_{shaft} = \Delta\rho = |\rho_2 - \rho_1|$, and for the distal long-bones ω_{shaft} was the width of the wider bone (as described in Section 6.2.2) and was calculated using $\omega_{shaft} = max(\triangle\rho_1, \triangle\rho_2)$.

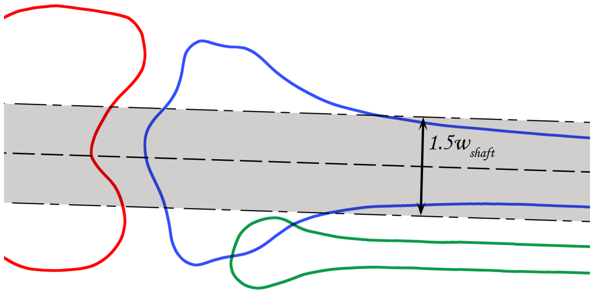

Figure 6.9: The search region of width $1.5\omega_{shaft}$ *around the bone centre-line.*

The value $1.5\omega_{shaft}$ was chosen for the width of the search region, that is, $0.75\omega_{shaft}$ on either side of the centre-line. This region was large enough to always contain the extreme articular surface—since in most cases this was normally located well inside ω_{shaft}—but not so large that other anatomical structures began to interfere. In addition, because this value was not based on the image size, but rather varied with the size of the bone being examined, it was a good estimate of the size of the required search region, regardless of how large the bone was within the image. The region of width $1.5\omega_{shaft}$ was extracted from both the smoothed and gradient magnitude images, as shown in Figure 6.10a and b.

The modified Hough Transform

To locate the tangent to the articular surface, the same modified Hough Transform discussed in Section 5.2.3 was used to calculate a single projection of the extracted gradient (Figure 6.10b) in the direction orthogonal to the bone centre-line. That is, rather than calculating the transform for a large number of θ in the range $[0, 180°)$, only the transform for the angle $\theta_{tangent}$ was calculated. The same value of $r = 2.5°$ in Equation 5.4 was used to match the direction of the gradient to the direction of the projection. This ensured that only edges almost orthogonal to the centre-line—such as those belonging to the articular surfaces—were included in the projection. The result was a plot that was close to zero for most values of ρ, but contained sharp peaks at those locations where there were strong edges orthogonal to the centre-line. In addition, a single projection of the extracted smoothed image (Figure 6.10a) in the same direction was calculated using the standard Hough Transform (that is, no angle

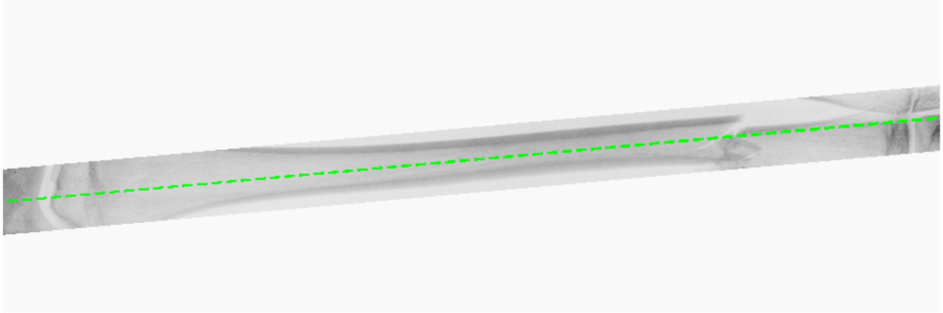

(a)

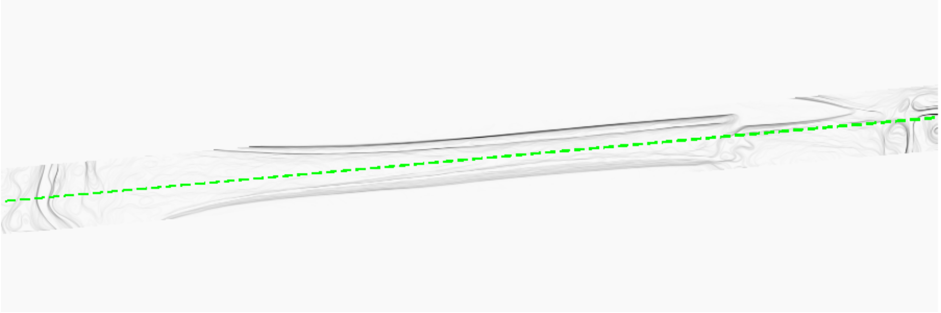

(b)

Figure 6.10: A bone segment of width $1.5\omega_{shaft}$ *located around the bone centre-line—corresponding to the lightly shaded region in Figure 6.9—was extracted from (a) the smoothed image, and (b) the gradient of the smoothed image.*

matching was performed). Different information was revealed by this second transform, because the peaks corresponded to dispersed areas of high intensity rather than edges in the direction of the projection. This projection therefore enabled efficient detection of the synovial cavity described in Section 2.2.1 on page 9. Characteristic troughs were created by the synovial cavity because the hyaline cartilage and synovial fluid were relatively radiolucent compared to the bone, so the value of the projection was much lower at that point. The resulting projections are shown in Figure 6.11.

Interpreting the resulting projections

As demonstrated in Figure 6.10b, there were unfortunately still a large number of lines present in the image that were orthogonal to the bone centre-line. This was exacerbated further if the subject was skeletally immature and the growth plates had not yet fused, resulting in an increased number of edges parallel to the extreme articular surface. These parallel lines manifested as large peaks in the gradient projection. Additional processing therefore had to be performed to accurately determine the true location of the articular surface. Combining the information from these two Hough Transform

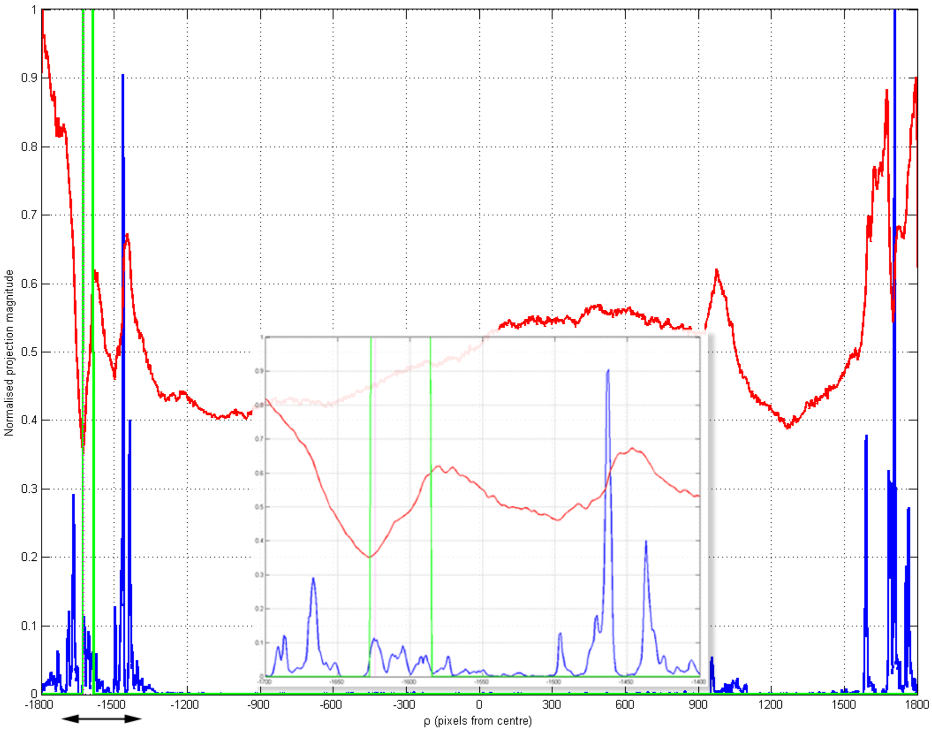

Figure 6.11: The projections of the extracted images shown in Figure 6.10a and b, in the direction orthogonal to the centre-line. The blue line is the projection of the gradient, calculated using the modified Hough Transform with $r = 2.5°$*, while the red line is the projection of the smoothed image, calculated using the standard Hough Transform. The green line is the articular surface search region, and the pink dashed line represents the location of the detected articular surface. The inset corresponds to the part of the x-axis with the arrow.*

projections (shown in Figure 6.11) allowed the extreme articular surface to be identified, since neither projection was sufficient by itself.

The synovial cavity was characterised by a large local (and in many cases, global) minimum in the projection of the smoothed image. Surrounding this point there were a number of local maxima in the gradient projection that corresponded to, among other things, the proximal and distal articular surfaces and the growth plates (if the subject was skeletally immature). To detect the synovial cavity, the smoothed image projection was searched to identify this local minimum. Once the synovial cavity was located, the articular surfaces had to be located within a small distance on either side of this point. Again, this distance was dependent on the shaft width, and was empirically estimated as being a maximum of $0.25\omega_{shaft}$. It was not necessary to search both sides of this point because the only articular surface of interest was the one that belonged to the bone whose centre-line had already been located. However, the correct side had to be identified before examination.

If the local minimum (and therefore the joint of interest) was in the left half of the image then it was assumed that the diaphysis to be examined was located to the right of it. Certainly this would be the case if sufficiently large peaks were to be obtained in the Hough accumulator such that the bone approximation parameters in Chapter 5 could be correctly determined. As a result, a region of width $0.25\omega_{shaft}$ to the right of the identified point was examined. The reverse was true when the identified local minimum was in the right half of the image.

The largest gradient peak in the region of width $0.25\omega_{shaft}$ that extended from the local minimum was identified. This point was the furthest tangent to an edge belonging to that bone, and so was marked as being the extreme articular surface. The green pulse shown in Figure 6.11 has a width $0.25\omega_{shaft}$, and its left side is located at the local minimum that corresponds to the synovial cavity. The largest gradient within the region marked by this pulse was chosen as the extreme articular surface. As shown in Figure 6.1, the articular surface could then be used as a reference point from which the distance to the diaphyseal boundary was measured. The identified articular surface is shown in Figure 6.12.

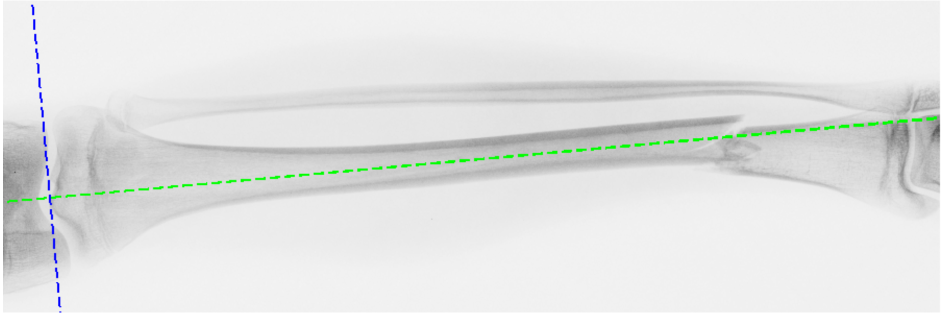

Figure 6.12: The green dashed line marks the detected long-bone centre-line, and the blue dashed line marks the orthogonal extreme articular surface.

6.3.2 Calculating the bounding box size

After the first and second segmentation steps from Figure 6.1 had been performed, the parameters for the centre-line(s) were known (Figure 6.1a), and the extreme articular surface had been located (Figure 6.1b). The next step was to determine the size of the bounding box shown in Figure 2.10. That is, how far away from the extreme articular surface the diaphyseal boundary was located. According to the AO definition, this distance was equivalent to the maximum width of the epiphysis $\omega_{epiphysis}$, as marked in Figure 6.1c. Importantly, as shown in this figure, the epiphyseal width that should be measured was that of the bone whose centre-line had been determined, not the width of the other bone in the joint of interest. This was critical, because in almost all cases the epiphyseal widths of the bones in the joint would be different, and measuring the wrong one would result in an incorrect segmentation.

Determining the epiphyseal width

The epiphyseal width was determined by summing the maximum deviation of the bone away from the centre-line on each side of the epiphysis. Tests showed that in most cases the bones were not symmetric, so the maximum deviations from the centre-line generally occured at different distances from the extreme articular surface tangent. Of course the maximum deviation had to occur within the bounding box defined in this step, but for all joints other than the proximal radius and ulna and proximal femur (which is already an exception to the bounding box rule shown in Figure 2.10), the maximum deviation occured very close to the articular reference line. It was found that a good estimate of this maximum distance was the value $1.5\omega_{shaft}$ that was utilised in the previous section. It was determined that lower values often produced sections that did not contain the points of maximum deviation, while higher values produced sections that could include the proximal fibula, resulting in a larger width measurement. As mentioned, it was necessary to ensure that when determining the epiphyseal width, that within the joint the correct bone was examined. That is, that the correct side of the articular surface tangent was analysed. The same method used to interpret the projections in the previous section was also used here. This region of width $1.5\omega_{shaft}$, measured from the extreme articular surface tangent, in which the maximum deviations from the centre-line must occur, is shown in Figure 6.13.

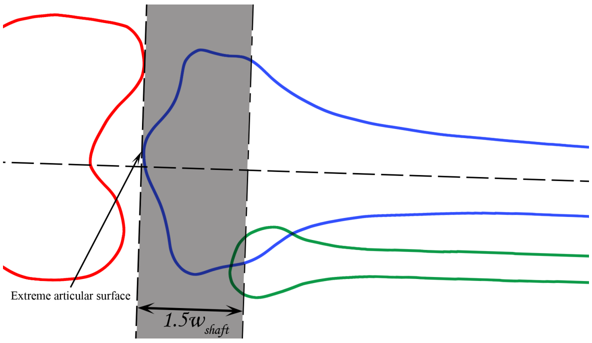

Figure 6.13: The search region of width $1.5\omega_{shaft}$, measured from the extreme articular surface.

A small region of the joint starting at the articular surface tangent and finishing at a distance $1.5w_{shaft}$ from that tangent was therefore extracted from the gradient magnitude image. An example of the extracted region is shown in Figure 6.14. To determine the epiphyseal width, a projection of the extracted sections in the direction of the centre-line was calculated using the modified Hough Transform, again using the same value of $r = 2.5°$ in Equation 5.4. Once again, this ensured that only edges that were almost parallel to the centre-line—such as those belonging to the epiphyseal edges—were included in the projection. The result was a plot that was close to zero for most values of ρ outside the bone, oscillated within the bone, and contained sharp peaks at those locations where there were strong edges parallel to the centre-line. The result is shown in Figure 6.15.

Since only edges parallel to the centre-line could contribute to the accumulator, an ideal projection would contain two large peaks that correspond to the two epiphyseal edges, and the identified bone centre-line would pass somewhere between the two

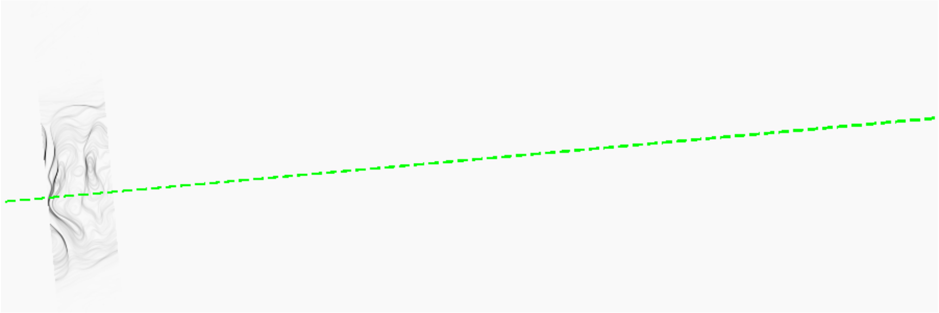

Figure 6.14: A bone segment of width $1.5\omega_{shaft}$ located at the extreme articular surface—corresponding to the lightly shaded region in Figure 6.13—was extracted from the gradient of the smoothed image.

peaks. Unfortunately, as illustrated in Figure 6.15, the actual projection was much more complicated, so the largest peaks on either side of the centre-line were selected as the epiphyseal edges. The epiphyseal width $\omega_{epiphysis}$ was the distance between these chosen maxima.

6.3.3 Diametaphysis identification

Once the location of the extreme articular surface and the epiphyseal width $\omega_{epiphysis}$ had been determined (the first three steps shown in Figure 6.1), the remaining side of the square marks the diametaphysis, the boundary between the metaphysis and diaphysis. This was located parallel to the extreme articular surface tangent, at a distance $\omega_{epiphysis}$ away from it on the same side as the search region. Depending on how much of the long-bone was visible in the x-ray image, this procedure could be performed again to determine the location of the diaphyseal boundary at the opposite end of the bone. In either case, the segment of the image corresponding to the diaphysis was then easily extracted. Figure 6.16a shows the construction lines laid over the original image, 6.16b shows the extracted diaphysis, and 6.16c shows the segmented region after the process was applied to the opposite end of the bone. It was in this region that the fracture detection algorithm described in Chapter 7 was applied. When the AO segmentation algorithm was applied to the six images in the development set, seven of the 9 required diametaphyses were correctly identified.

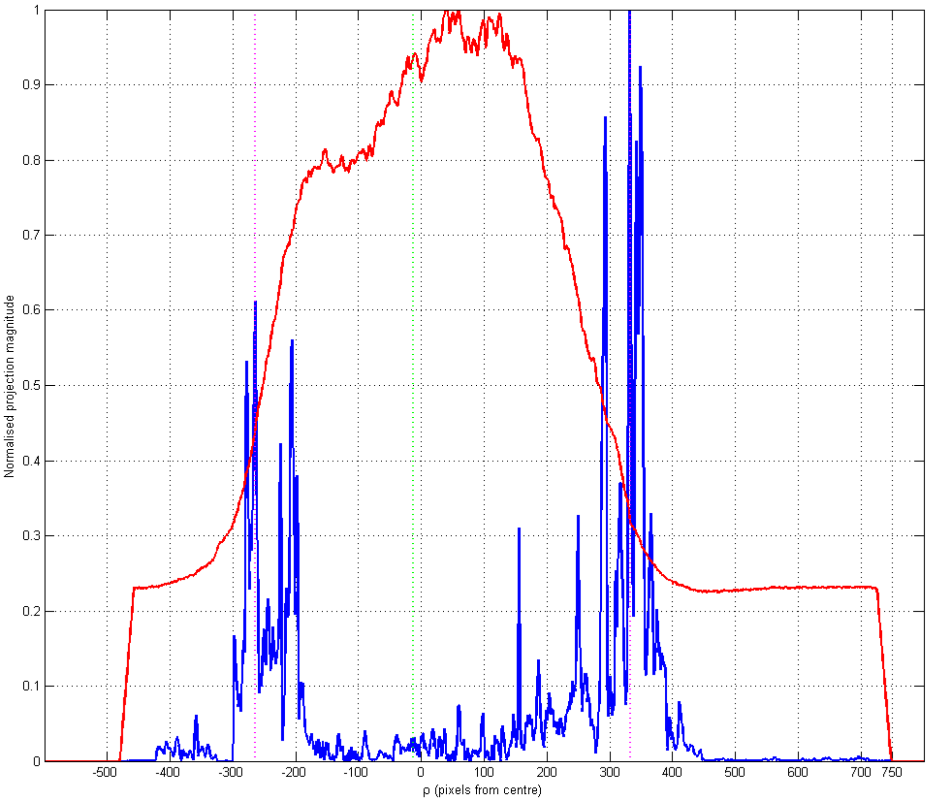

Figure 6.15: The projection of the gradient of the extracted region shown in Figure 6.16b is shown in blue. The green dashed line marks the location of the bone centre-line, and the pink dashed lines mark the maxima on either side of the centre-line that corresponded to the edges of the epiphysis. The epiphyseal width $\omega_{epiphysis}$ *is the distance between these two pink dashed lines. The projection of the smoothed image is shown in red, but was not used for calculating the epiphyseal width.*

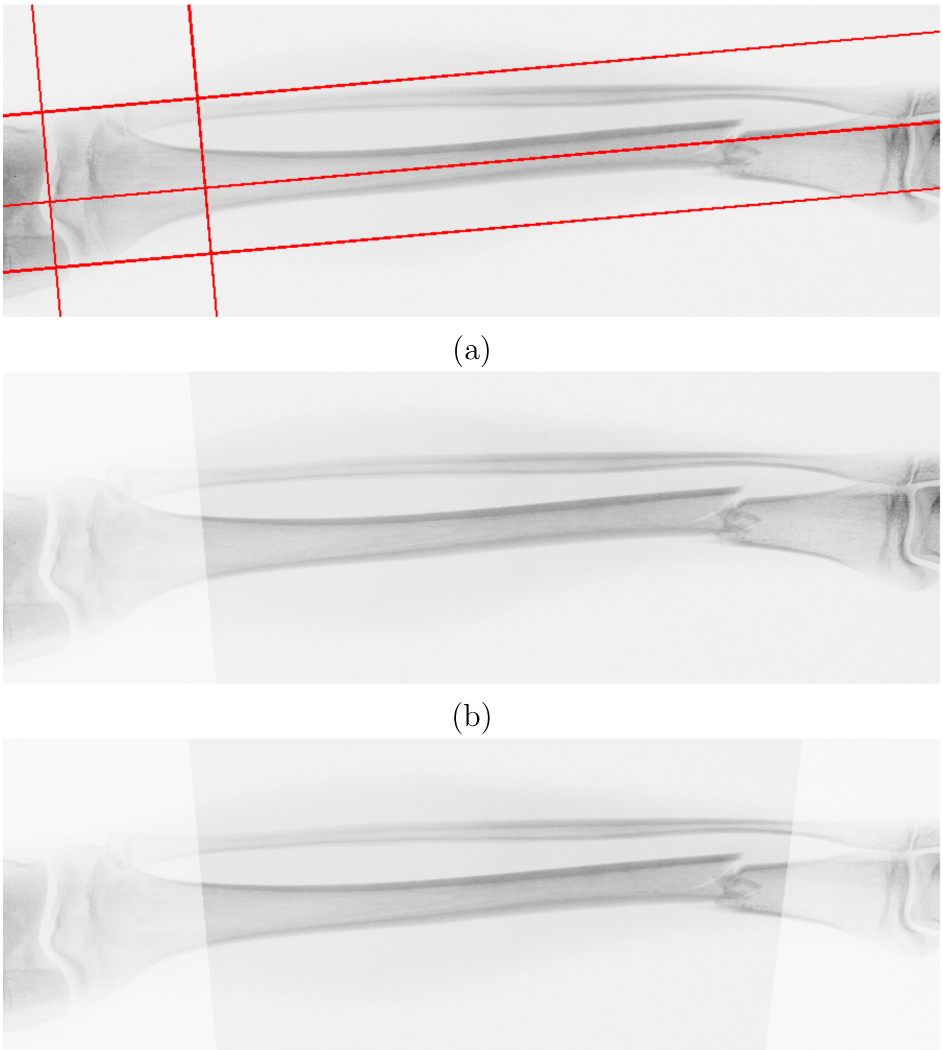

(a)

(b)

(c)

Figure 6.16: The completed segmentation, showing (a) all the construction lines used for the proximal segmentation, and (b) the partially segmented diaphysis, and (c) the completed segmentation after the method was applied a second time to locate the distal diaphyseal boundary.

6.4 Segmentation using the bone curvature

Although capable of segmenting many long-bone diaphyses, the AO segmentation method discussed so far in this chapter suffered from a few drawbacks that limited its effectiveness. Four limitations of the AO segmentation method are outlined below:

1. While the centre-line estimation algorithm produced very good results, the articular surface identification and epiphyseal width calculation algorithms were both prone to detecting the wrong features within the image. Some reasons why the wrong features were detected are discussed in Section 6.6. Obviously, incorrect detection of either the articular surface or the epiphyseal edges would result in an inaccurate segmentation.

2. As previously mentioned at the start of Chapter 6, the AO segmentation method was only applicable to AP images, since the epiphyseal width can only be measured in an anteroposterior view. The complete 44 image test set contained 29 AP images (66%) and 15 ML images (34%), so the AO segmentation method could only be applied to two thirds of the image set, leaving one third of the images unable to be segmented due to the x-ray view.

3. The AO segmentation method also required the epiphysis to be completely contained within the image, so that both the articular surface and epiphyseal edges could be identified, and all required measurements could be made. It was not always the case that these three landmarks were contained within the image. For example, neither end of the femur shown in Figure 6.4 was visible, so the AO segmentation could not be applied. Of the 29 AP images in the test set, 7 (24%) did not contain sufficient landmarks to allow any AO segmentation. This means that taking into account the requirements of both an AP view and complete epiphysis inclusion, only 22 (50%) of the images in the complete test set could have any AO segmentation applied to them.

4. As shown in Figure 2.10 on page 18, the AO segmentation method did not specify the point at which the diaphysis began at the proximal end of the femur. Three

of these images were present in the test set, reducing the number of images that could be adequately segmented using the AO method to 21 (47%).

In addition, the accuracy with which the diaphysis segmentation conformed to the AO method was not critical. This was because every bone was different, and the AO method was only a guide as to where the diametaphyses were likely to be located. To overcome these significant limitations, a better method of performing the segmentation was required.

6.4.1 Diametaphysis identification

Although titled diaphysis segmentation, the aim of this chapter—and the AO segmentation method already described—was essentially to locate the diametaphysis, the region, as stated in Section 2.1.1, in which the long-bone shaft begins to widen and curve. Fortunately, the long-bone shaft endpoint detector examined in Section 5.7.3 was used to locate those points at which the bone deviated sufficiently away from the parameterised line fitted to the straight portion of the bone edge. By their nature, the points at which the deviation occured were the diametaphyses. To enable a robust segmentation in all cases, these points were utilised, along with some features from the previously outlined AO segmentation.

Since the line endpoints on either side of the bone often varied, in some cases to quite a large extent, it was appropriate to use an average measure. The average location had already been calculated using Equation 6.2, and could be utilised again here. To ensure that the correct point on the diaphysis was selected, the most appropriate centre-lines were still detected as described in Section 6.2.2. As with the AO method, the segmentation boundary at the diametaphysis should be orthogonal to the centre-line at that point. Therefore at both ends of every identified centre-line, an orthogonal line spanning the image was added, as shown in Figure 6.17. For these lines, the θ parameter could be calculated because the angle of the centre-line was known, and the ρ parameter could be calculated using Equation 5.1 because the line endpoints (x, y) were also known. Some of these lines would eventually become the locations at which the segmentation was performed.

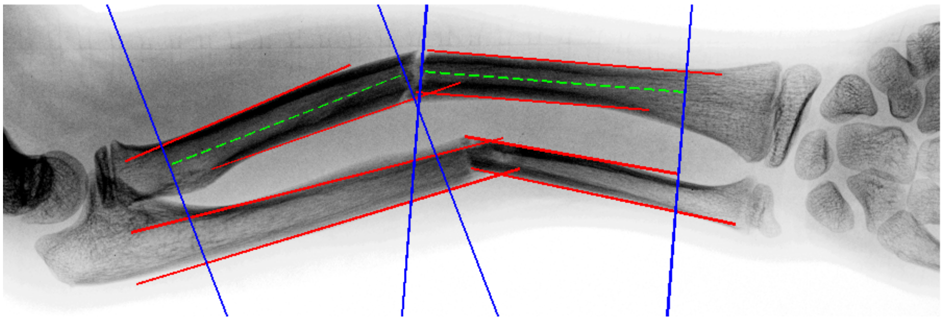

Figure 6.17: Once the appropriate centre-lines (green dashed) had been located from all of the bone edges (red solid), lines orthogonal to the centre-line endpoints (blue solid) were added. From these lines the segmentation was performed.

In some cases—such as undisplaced proximal fractures where there was only a single centre-line—the segmentation was performed at the two identified lines. In more complex cases, such as the image shown in Figure 6.17, it was necessary to select two lines from a larger set. Regardless of the number of centre-lines in the image, it was always the two orthogonal lines furthest apart that corresponded to the diametaphyses. These could be selected by ranking the lines based on the ρ parameter, and then selecting the first and last parameter pairs as the points at which the segmentation would be performed.

When this method was applied to the six development images, all six were segmented correctly. Figure 6.18 shows the identified diametaphyses in four of the development images. In both 6.18a and 6.18b the epiphyses were not present in the image (the third limitation identified in this section), but unlike with the AO segmentation, the diaphyses were accurately located based on the centre-line endpoints. Comparing the same development image segmented using the AO method (Figure 6.16c) and the bone curvature method (Figure 6.18d) showed that the two methods produced similar results. For this particular image, the only significant difference was that the segmentation line on the right side of the bone was oriented at a different angle, because the centre-line of the short fragment was not detected in the previous stage. However, this did not affect the subsequent stages significantly. Again, once the segmentation was performed, the fracture detection algorithm described in Chapter 7 could be applied to the segmented region.

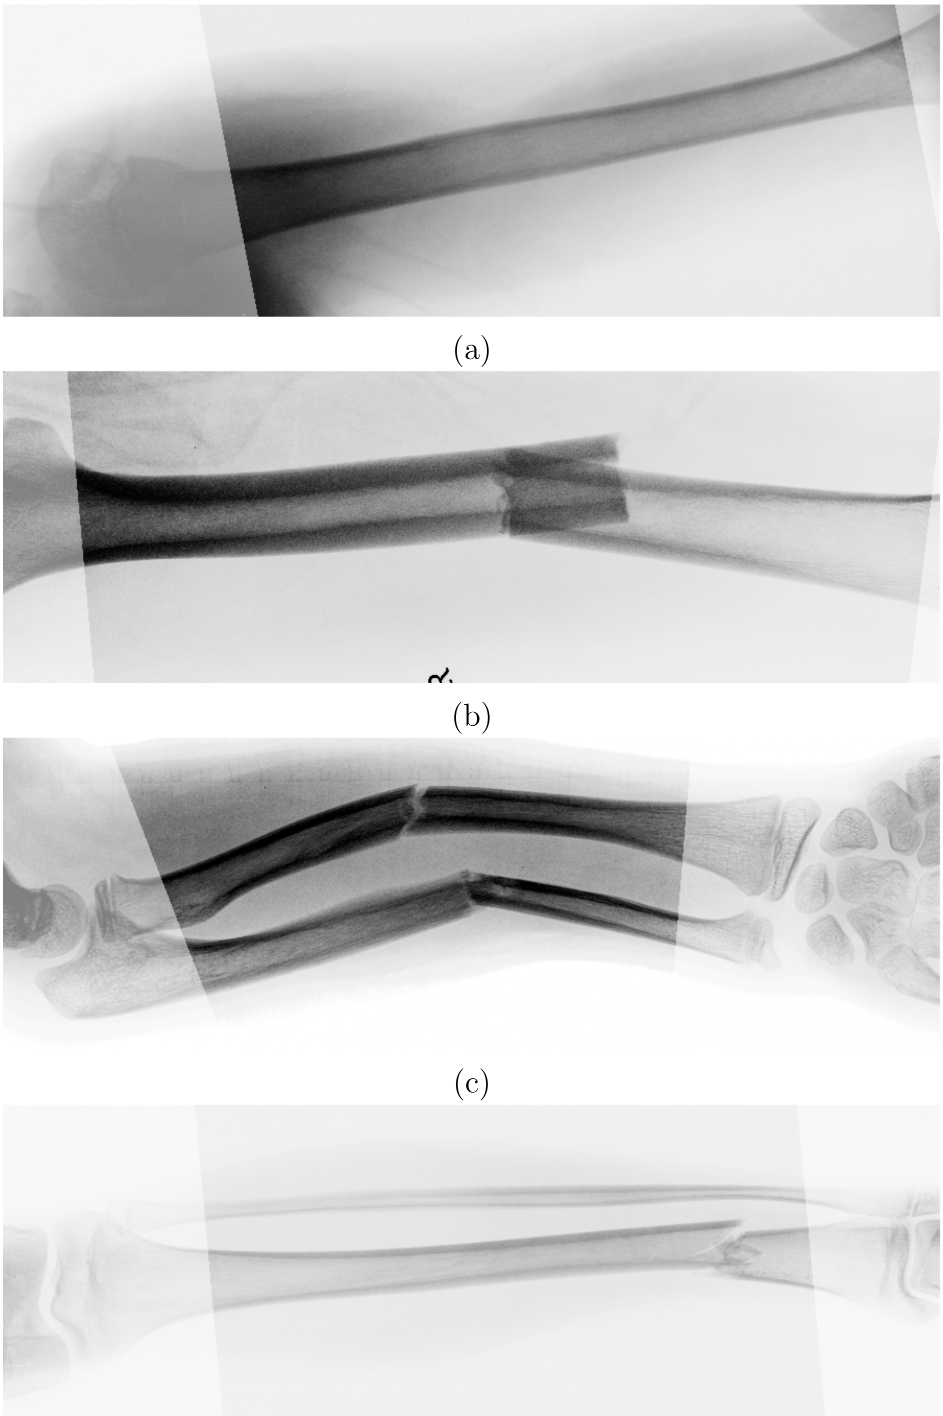

(a)

(b)

(c)

(d)

Figure 6.18: Examples of four development images showing the accurate identification of the diametaphysis, based on the centre-line endpoints.

6.5 Increasing segmentation accuracy

The accuracy with which the segmentation was performed could be improved by utilising some user intervention. Although this meant that the segmentation was no longer fully automatic, in some cases it improved accuracy, especially in very complicated images. The algorithm was constructed to allow the user to intervene when the segmentation was poor, as could occur when the bone landmarks were not correctly identified. Once the segmentation was complete, the user was able to manually move the boundaries to correctly match the diametaphysis, if this intervention was required.

6.6 Diaphysis segmentation evaluation

The two algorithms developed to perform long-bone diaphysis segmentation were evaluated by applying them to the images in the test set. For all images, the results obtained from the computer algorithms were compared to segmentations performed manually and were deemed to be correct if the segmentation points were located within a distance of $0.25\omega_{shaft}$ on either side of the manually selected point. Additionally, for the bone centre-line estimation, a centre-line was deemed to be correct if it was a good visual approximation.

Bone centre-line estimation

When tested, the centre-line detection algorithm correctly located 52 of the 56 centre-lines (93%) that were manually identified in the 44 images. In all four cases where a centre-line was not detected correctly, the cause was that the lines were not paired appropriately by the algorithm. An example of one of these images is shown in Figure 6.19. This incorrect pairing occurred because the medial edge of the radius and the lateral edge of the ulna were paired, because they were in such close proximity. This also meant that in these four cases, the required combination did not necessarily produce the minimum distance between the line endpoints. Fortunately in these cases, the incorrect detection of the centre-line had little effect on the final outcome of the segmentations produced by either of the algorithms, although this would not always be the case. In

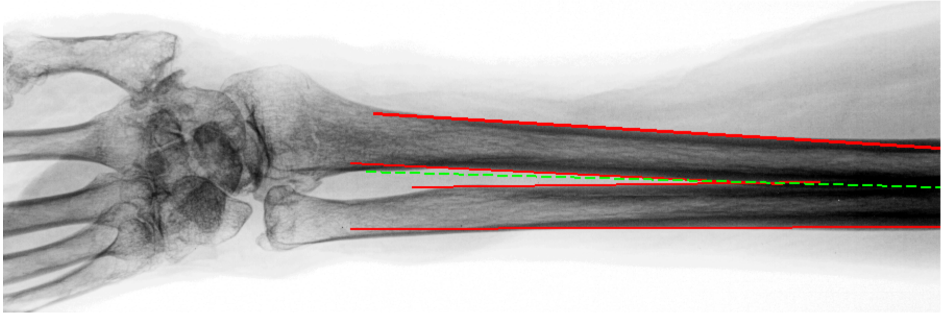

Figure 6.19: An example of incorrect centre-line detection that occurred because pairing the two internal lines, and the two external lines resulted in the smallest total distance between the line endpoints.

all cases where the complete set of centre-lines were correctly detected, the selection of the most appropriate centre-lines was performed with 100% accuracy.

Segmentation using the AO method

After the centre-lines were detected, the segmentation algorithm utilising the AO method was evaluated. As discussed at the start of Section 6.4, the AO method had a number of inherent limitations that reduced the number of images in the test set to which the algorithm could be applied, from 44 down to 21. This severe limitation implied that less than half of the images in the test set (and therefore a standard sample of images likely to be viewed in an emergency department) were suitable for segmentation by the AO method. In addition, when the algorithm was applied to those 21 images, only six of the 30 segmentation boundaries (20%) were detected correctly.

In almost all cases ($> 90\%$), the extremely poor performance of the algorithm was due to the lack of robustness of the articular surface identification step. There was a huge natural variation between images that made an appropriate articular surface identification criteria (for use in the projection of the segment) almost impossible to determine. The results did show that this step was much more likely to succeed when the synovial cavity was large and the bones were completely separated in the x-ray projection. In most cases, bones overlapping within the joints caused the algorithm to fail because the articular surface could not be identified. In addition, the presence of growth plates (Section 2.2 on page 9) also affected the segmentations of a number of images, since the physis could easily be mistaken for the articular surface.

Evaluation of the algorithm with a different set of criteria for identifying the articular surface was also performed. Utilising the gradient of the projection as the primary criterion resulted in a different set of images correctly segmented. However, in all cases the total number correctly segmented was never greater than 20%. This led to the conclusion that locating the articular surface using only the projection data was a flawed solution, and that accordingly the AO segmentation method was not a suitable approach for semi-automated long-bone diaphysis segmentation.

Segmentation using the bone curvature method

The segmentation algorithm that utilised the bone curvature was also evaluated using the same criteria, to determine its accuracy, and to compare it to the AO segmentation method. This time, the algorithm could be applied to all 44 images, rather than a small subset as was the case with the AO segmentation method. Of the 71 segmentation boundaries required in the image test set, 59 were identified correctly, corresponding to a detection rate of 83%. In addition, in 33 out of the 44 images (75%), all the required segmentation boundaries were correctly detected, such that those images were segmented completely correctly.

In six images the algorithm detected an extra segmentation point, where one was not required. This corresponded to 8.5% of the total number of lines. In almost all cases the extra segmentation point was close to the edge of the image because the centre-line did not extend all the way to the edge of the image, as demonstrated in Figure 6.20. In all six cases the extra segmentation point was easily removed by user intervention. The most common cause of choosing the incorrect location for segmentation (6 cases) was the fracture disturbing the centre-line such that it ended prematurely, rather than continuing to the point where it would if the fracture were not present. Thus, segmentation using the bone curvature was better in images where the bones were complete—that is, images containing subtle fractures with no displacement. In fact, when a subset of 24 images containing only subtle undisplaced fractures was used—rather than the complete set that also contained obvious grossly displaced fractures—the diaphysis segmentation rate was 100%. Fortunately the long-bone segmentation and fracture detection algorithm is aimed at these types of images, since it is subtle fractures that

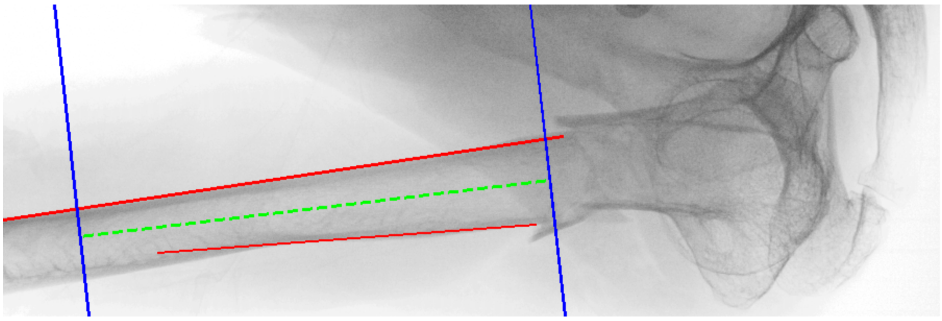

(a)

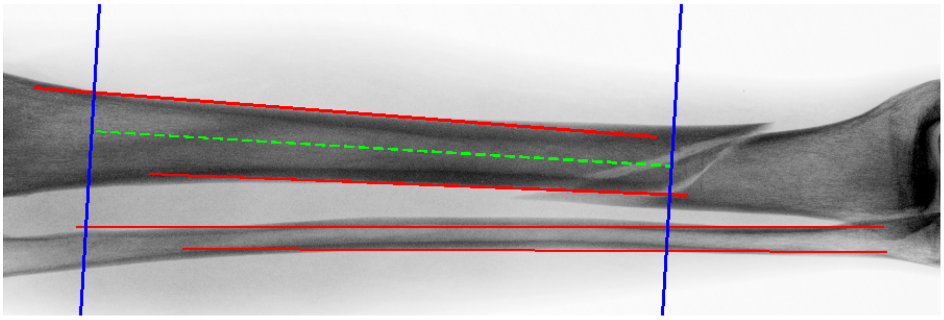

(b)

Figure 6.20: Examples showing incorrect segmentations produced by the bone curvature method. (a) A segmentation point where one was not required (blue line on the left of the image) was caused by a short centre-line (green dashed line) that didn't extend to the edge of the image. (b) An incorrect segmentation location (blue line on the right of the image) was caused by the fracture disrupting the bone centre-line (green dashed line).

are most likely to be missed during reading.

Clearly the segmentation method that utilised the bone curvature—via the centre-line endpoints—was far superior to the AO method, not only in terms of accuracy but also in terms of the size of the set of images to which it could be applied. As such, it was the method of choice for semi-automatic long-bone segmentation.

6.7 Summary

This chapter described how the approximated long-bone shaft parameters (ρ, θ) calculated in Chapter 5 could be used to perform long-bone diaphysis segmentation using two different methods. The first of these was a segmentation algorithm based on the AO scheme described in Section 2.3.4, that utilised four distinct stages. The first stage involved locating the bone centre-line, or centre-lines depending on the image. Three centre-line identification methods were described, and their limitations noted. The final scheme—utilised for long-bone segmentation—involved choosing the combination of pairs that minimised both the total distance between the line endpoints, and the total distance between the lines, calculated using the ρ values. This method was the most accurate with the development images, with all 28 lines correctly paired to produce 14 centre-lines.

In the second stage, after the centre-lines were accurately located, it was necessary to identify the articular surface. The point of interest—called the extreme articular surface—was described as the point on which the bone would rest if it were held upright on a flat surface. This point was normal to the centre-line previously calculated, so that locating the extreme articular surface was equivalent to locating the furthest tangent to the bone edge that was perpendicular to that centre-line. To simplify this task, a search region of width $1.5\omega_{shaft}$ around the centre-line was used. Within that search region, the modified Hough Transform was used to detect lines orthogonal to the centre-line. The resulting projections were then interpreted to ensure that the correct location—the articular surface adjacent to the synovial cavity—was identified.

Once the centre-line and extreme articular surface were identified, the third stage required that the bounding box size be determined. This required that the epiphyseal width be accurately measured, so that the boundary between the diaphysis and metaphysis—the diametaphysis—could be located. This again required a search region of width $1.5\omega_{shaft}$ from the articular surface, in which the modified Hough Transform was used to locate lines parallel to the centre-line. Once this measurement was made, stage four utilised this width to mark the diametaphysis. Diaphysis segmentation was performed correctly for seven of the nine required points in the development image set.

Unfortunately this implementation of the AO scheme suffered from a number of limitations that made it inappropriate for diaphysis segmentation. Hence, a second method based on analysis of the bone curvature was formulated. In this alternative algorithm, the ends of the centre-lines were used to determine the segmentation points, since they corresponded to the average of the two line endpoints from the parameter approximation. The line endpoints could be used because they are the points at which the bone began to curve and deviate away from the straight line approximations. By definition these points were the diametaphyses. In cases where multiple centre-lines were present, the appropriate lines were selected by ranking them based on their ρ values, and then selecting the first and last parameter pairs. This method proved to be much more robust for accurate segmentation, with all six development images correctly segmented.

Results from testing the automated long-bone segmentation method showed that close to 95% of the required centre-lines were detected correctly, based on the long-bone shaft approximation parameters found in Chapter 5. In addition, in all cases where the complete set of centre-lines were correctly located, the appropriate centre-lines were selected from this set with 100% accuracy. The AO segmentation method proved to not only be limited in the types of images that it could be used to segment, but also limited by its accuracy. Only 20% of the small subset of 21 images contained one correct segmentation point. In contrast, the segmentation method that utilised the bone curvature was much more successful in locating the correct segmentation boundaries, with almost 85% of the full set of 71 points located correctly. In all other cases where the segmentation boundaries were not correct, they were easily modified by the user moving them to the correct locations.

Diaphysis segmentation needed to be performed before a fracture detection algorithm could be applied to the identified diaphysis. The next chapter describes a method of detecting fractures within the marked diaphysis.