Computing Neck-Shaft Angle of Femur for X-Ray Fracture Detection

Tai Peng Tian[1], Ying Chen[1], Wee Kheng Leow[1], Wynne Hsu[1], Tet Sen Howe[2], and Meng Ai Png[3]

[1] Dept. of Computer Science, National University of Singapore
3 Science Drive 2, Singapore 117543
{tiantaip, leowwk, whsu}@comp.nus.edu.sg
[2] Dept. of Orthopaedics, Singapore General Hospital
Outram Road, Singapore 169608
tshowe@yahoo.com
[3] Dept. of Diagnostic Radiology, Singapore General Hospital
Outram Road, Singapore 169608

Abstract. Worldwide, 30% – 40% of women and 13% of men suffer from osteoporotic fractures of the bone, particularly the older people. Doctors in the hospitals need to manually inspect a large number of x-ray images to identify the fracture cases. Automated detection of fractures in x-ray images can help to lower the workload of doctors by screening out the easy cases, leaving a small number of difficult cases and the second confirmation to the doctors to examine more closely. To our best knowledge, such a system does not exist as yet. This paper describes a method of measuring the neck-shaft angle of the femur, which is one of the main diagnostic rules that doctors use to determine whether a fracture is present at the femur. Experimental tests performed on test images confirm that the method is accurate in measuring neck-shaft angle and detecting certain types of femur fractures.

1 Introduction

Many people suffer from fractures of the bone, particularly the elderly folks. According to the findings of the International Osteoporosis Foundation [1], the lifetime risk for osteoporotic fractures in women is 30% – 40% worldwide, and 13% in men. The number of hip fractures could rise from 1.7 million worldwide in 1990 to 6.3 million by 2050. Most dramatic increase is expected to occur in Asia during the next decades. According to World Health Organization, osteoporosis is second only to cardiovascular disease as a leading health care problem [1].

In practice, doctors and radiologists in the hospitals rely mainly on x-ray images to determine whether a fracture has occured and the precise nature of the fracture. Manual inspection of x-rays for fractures is a tedious and time consuming process. Typically, the number of images containing fractures constitute a small percentage of the total number of images that the radiologists have to scan through. For example, in our test images, only 11% of the femurs are fractured. After looking through many images containing healthy femurs, a tired

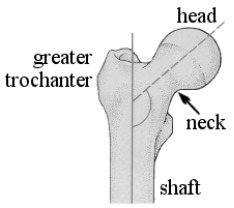

Fig. 1. Neck-shaft angle is the angle made by the shaft axis (solid line) and the neck axis (dashed line) of the femur.

radiologist has been found to miss a fractured femur among the many healthy ones. As some fractures are easier to identify than others, an automated fracture detection system can assist the doctors by performing the first examination to screen out the easy cases, leaving a small number of difficult cases and the second confirmation to the doctors. Automated screening of both healthy and fractured cases can thus relieve some of the labor intensive work of the doctors and help to improve the accuracy of their diagnosis. Therefore, this computer vision application is extremely useful for clinicians and is now feasible because all clinical radiology is going digital. Digital x-ray images are now routinely captured using digital x-ray machines.

Among the various fracture incidents, common hip fractures of the femur account for the largest proportion of fracture cases. One of the main diagnostic rules that doctors use to detect femur fracture is by assessing the distortion of the so-called *neck-shaft angle*, that is, the angle between the shaft and the neck axes (Fig. 1). The neck-shaft angle of a healthy adult femur is about 120 to 130 degrees. A large discrepancy from the healthy neck-shaft angle would indicate a possibility of fracture. Thus, this article focuses on the automated measurement of neck-shaft angle of the femur in an x-ray image and uses the measured angle to determine whether a fracture has occurred.

At first glance, it may seem that automated measurement of neck-shaft angle is a trivial task for a computer. However, it turns out to be far from trivial, especially for fractured femurs. The femoral neck usually appears as a very short segment in an x-ray image. Correct localization of the neck is a very difficult task. For certain types of fractured femurs, the necks are crashed and do not even appear on the x-ray images (Fig. 2(c, d)). To overcome this problem, we define the neck axis to be the axis of symmetry of the 2D contour of the femoral head and neck, and applies an optimization algorithm to determine the best fitting symmetry axis (Section 3).

2 Related Work

So far, we have not come across any published work on the computer automated detection of fractures in x-ray images. The closest related methods used non-visual methods to detect fractures. For example, Ryder et al. analyzed acoustic pulses as they travel along the bone to determine whether a fracture has oc-

curred [2]. Kaufman et al. applied a neural network model to analyze mechanical vibration [3] whereas Singh and Chauhan measured electrical conductivity [4].

Most of the research efforts related to orthopaedics have instead been focused on the detection of osteoporosis (e.g., [5–7]). These methods of detecting osteoporosis usually assume that an area of interest is provided by the operator. So, there is no need to automatically detect the contour of the bones under examination. The image analysis required for the detection of osteoporosis is, therefore, simpler than that for fracture detection.

There are substantial amount of work on the analysis of tubular structures such as blood vessels and lung bronchi. In the analysis of these small structures, it is reasonable to assume certain relationship between image intensity and the position on the structure. For example, the cores method [8] and the ridge detection method [9] have been applied to 2D images to find intensity ridges which correspond to the medial lines of vessels. However, the femur is a large structure with complicated internal structure, which shows up as complex texture patterns in an x-ray image. So, standard method of analyzing tubular structures cannot work on the x-ray images of femurs.

3 Fracture Detection Method

Our method of detecting fractures in the femur consists of three stages: (1) extraction of femur contour, (2) measurement of neck-shaft angle, and (3) classification of femur based on measured neck-shaft angle. The extraction of femur contour in stage 1 is performed using a combination of algorithms, namely Canny edge detection and Hough transform for detecting significant straight line and curve features, and active contour mode (i.e., elastic snake) [10] with Gradident Vector Flow (GVF) method [11] to snap on to the continuous femur contour based on the line and curve features detected. Due to space limitation, this paper will focus on stage 2 (Section 3.1) of the process. Stage 3 will be discussed together with the experimental results (Section 4).

3.1 Measurement of Neck-Shaft Angle

To measure the neck-shaft angle, we have to recover the shaft axis and the neck axis. However, standard algorithms, such as the medial axis transform, are too sensitive to the noise along the contour, and they fail completely to extract the neck axis especially for fractured femur where the neck is crushed and distorted. Instead, a more robust algorithm that exploits the shape of the femur is used.

Computing the Shaft Axis. Note that the contour lines along the femoral shaft are almost parallel. If lines normal to the shaft contours are drawn from one side of the shaft to the opposite side, then the mid-points of the normal lines would be aligned with the shaft axis. We call these normal lines *level lines* as each line denotes a level along the shaft. (Note that our "level lines" are different from the well-known "level set" algorithm.) It turns out that level lines can also be

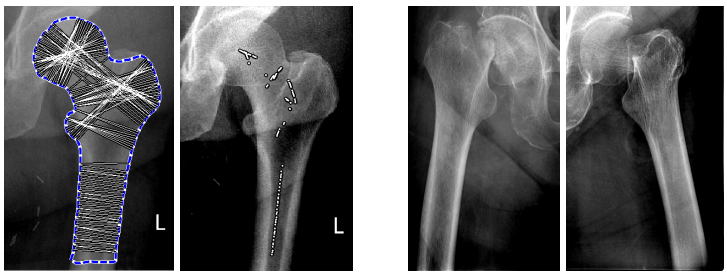

(a) (b) (c) (d)

Fig. 2. (a) Level lines found in the femur contour. (b) Mid-points of the level lines at the shaft are oriented along the shaft axis. (c, d) The neck contours are compressed or absent in these fractured femurs.

found eminating from the femoral head, passing through the approximate center of the head, and ending at the lower part of the greater trochanter. Whereas the level lines at the shaft are perpendicular to the shaft axis, those at the neck are parallel to the neck axis.

Instead of computing all possible level lines, we compute only those that intersect the snake points. First, the unit normal vector $\mathbf{n}_i$ of snake point $\mathbf{p}_i$ is computed by applying Principal Component Analysis (PCA) on a neighborhood of snake points centered at point $\mathbf{p}_i$. The first eigenvector of PCA would be tangential and the second eigenvector normal to the contour at point $\mathbf{p}_i$. Then, a line $l(\mathbf{p}_i, \mathbf{p}_j)$ joining points $\mathbf{p}_i$ an $\mathbf{p}_j$ is a level line if it is parallel to the normals $\mathbf{n}_i$ and $\mathbf{n}_j$, i.e.,

$$|\mathbf{n}_i \cdot \mathbf{n}_j| \approx \frac{|(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{n}_i|}{|\mathbf{p}_i - \mathbf{p}_j|} \approx \frac{|(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{n}_j|}{|\mathbf{p}_i - \mathbf{p}_j|} \approx 1 \ . \tag{1}$$

In the current implementation, two orientations $\mathbf{v}_1$ and $\mathbf{v}_2$ are considered similar i.e., $|\mathbf{v}_1 \cdot \mathbf{v}_2| \approx 1$ if $|\mathbf{v}_1 \cdot \mathbf{v}_2| \geq 0.98$.

Figure 2(a) shows an example of the level lines found in the femur contour. The level lines at the shaft can be easily isolated from the other level lines because they are short and are located at the lower half of the image. Given the shaft level lines, the mid-points of the level lines are computed (Fig. 2b) and a straight line is fitted through the mid-points to obtain the shaft axis.

Computing the Neck Axis. Figure 2(a) shows that there are several bundles of level lines within the femoral head and neck region. An adaptive clustering algorithm similar to that in [12] is applied to cluster the level lines into bundles according to two criteria: (1) the lengths of the lines, and (2) the mid-points of the lines. Level lines with similar lengths and whose mid-points are close to each other are clustered into a bundle. The algorithm is adaptive and it can compute the appropriate number of bundles required. After clustering, the bundle with the largest number of long level lines is chosen and the mean direction of the level lines is computed to approximate the orientation of the neck axis.

The above algorithm works well for healthy femur. However, for fractured femur whose neck is crushed, the contours of the neck may not even exist (Fig. 2(c, d)), complicating the problem of determining the neck axis. To obtain a more accurate estimation of the neck axis, an optimization algorithm is applied to compute the axis of symmetry of the femoral head and neck given the initial estimate obtained using the algorithm described above. Before computing the axis of symmetry, the femur contour is first smoothed with a 1-D Gaussian filter to remove noise along the contour. The σ value of the Gaussian should be large enough to produce smooth contour at the head and neck regions but not too large that the shape of the contour is severely distorted. In the current implementation, a σ value of 5 is used.

The general idea of computing axis of symmetry is to find a line through the femoral head and neck such that the contour of the head and neck coincides with its reflection about the line. Given a snake point $\mathbf{p}_k$ along the head contour, the mid-point $\mathbf{m}_i$ along the line joining snake points $\mathbf{p}_{k-i}$ and $\mathbf{p}_{k+i}$ is computed. That is, we obtain a midpoint for each pair of snake points on the opposite sides of $\mathbf{p}_k$. Then, a line l_k is fitted through the midpoints $\mathbf{m}_i$ to obtain a candidate axis of symmetry. If the contour is perfectly symmetrical, and the correct axis of symmetry is obtained, then each contour point $\mathbf{p}_{k-i}$ is exactly the mirror reflection of $\mathbf{p}_{k+i}$. So the fitting error E_k for l_k is

$$E_k = \frac{1}{n} \sum_{i=-n/2}^{n/2} \left| \mathbf{p}_{k+i} - \mathbf{p}'_{k-i} \right| \tag{2}$$

where $\mathbf{p}'_{k-i}$ is the reflection of $\mathbf{p}_{k-i}$ about l_k. E_k indicates how well is l_k an axis of symmetry. The best axis of symmetry is a line l_t that minimizes E_k. This procedure can be completed in $O(n^2)$ time where n is the number of snake points along the head and neck contour.

In the current implementation, the extent of the neck-head contour is empirically determined to be 40 snake points each to the left and right of a given snake point on the femoral head. The start position of the optimization algorithm is the snake point that is closest to the approximate neck axis computed in the previous stage using the level-line method. 20 snake-point positions to the left and right of the start position are considered in finding the best-fitting position and orientation of the neck axis.

4 Experimental Results

63 images each with a left and a right femur were used as the training images. The neck-shaft angles of these 126 femurs were computed and the decision threshold that minimized the number of detection error was determined (Fig. 3). This threshold was determined to be 116°, i.e., femurs with neck-shaft angles smaller than 116° were classified as fractured. This threshold value was used for classifying the other 160 test images containing a total of 320 femurs.

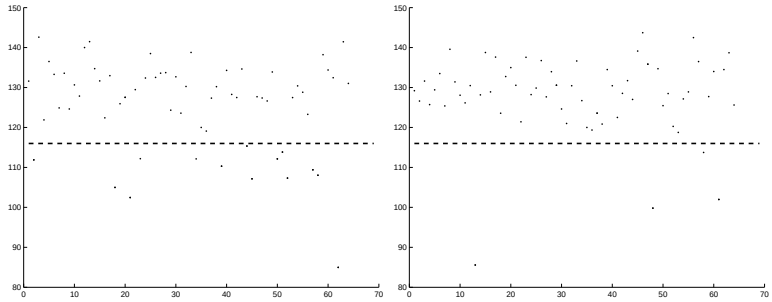

(a) (b)

Fig. 3. Measured neck-shaft angles of training images. (a) Left femurs. (b) Right femurs. Dots: healthy femurs, squares: fractured femurs, dashed line: decision threshold.

Table 1. Summary of training and testing results.

	training			testing		
classified	left femur	right femur	both	left femur	right femur	both
correctly as fractured	12	3	15	8	7	15
correctly as healthy	48	56	104	134	147	281
sub-total	60	59	119	142	154	296
	(95.2%)	(93.7%)	(94.4%)	(88.8%)	(96.3%)	(92.5%)
incorrectly as fractured	2	1	3	5	1	6
incorrectly as healthy	1	3	4	13	5	18
sub-total	3	4	7	18	6	24
	(4.8%)	(6.3%)	(5.6%)	(11.3%)	(3.8%)	(7.5%)
Total	63	63	126	160	160	320

Table 1 summarizes the classification performance on the training and testing data. 94.4% of the training images and 92.5% of the testing images are correctly classified. Figures 4 and 5 illustrate sample femurs that are correctly classified.

Visual inspection of the wrongly classified femurs indicate two main sources of error in the training and testing data. First, some fractured femurs are missed by the algorithm because there is no significant change of neck-shaft angle. These fractured femurs can be divided into two main categories:

1. Some fractures at the femoral necks cause the femoral heads to be displaced along the neck axes. As a result, there are no significant changes of neck-shaft angles though the shapes of the femoral head and neck regions are changed.
2. Some fractures are very slight cracks of the bones, and others are complete breakage of the femoral necks without significant displacements of the femoral heads. In these cases, there are no significant changes of both the neck-shaft angles and the shapes of the femoral head and neck regions.

These fractures can only be detected using other criteria and methods.

The second source of error is due to the misclassification of healthy femurs as fractured. These cases can be categorized into two types:

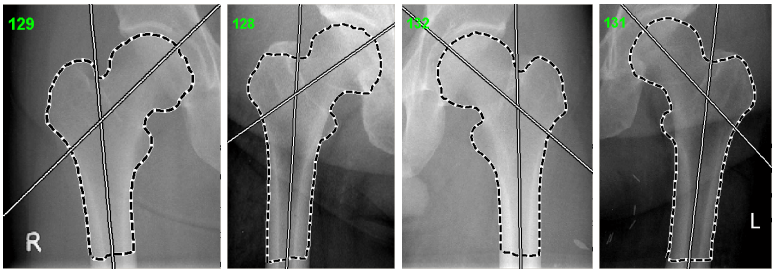

Fig. 4. Femurs correctly classified as healthy.

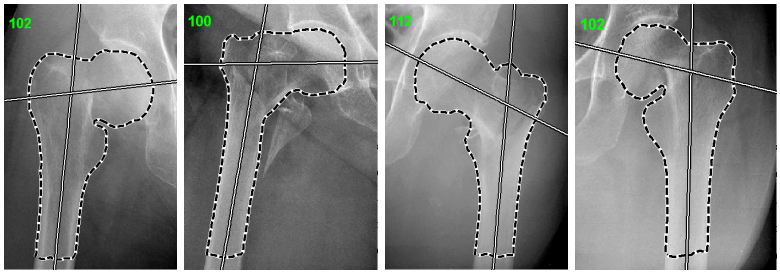

Fig. 5. Femurs correctly classified as fractured.

1. The femoral shafts are either very short or completely missing in the x-ray images. As a result, it is impossible to compute the shaft axes correctly.
2. The x-ray images are taken at unusual pose causing the shapes of the femurs to be distorted.

The first problem can be solved using methods that do not need the shaft to make correct classifications, and the second can be solved only with 3D models.

The above two sources of complications account for the majority of the misclassifications (Table 2). Only one femur among the training images and two among the testing images are wrongly classified due to the failure of the algorithm in computing the correct neck-shaft angle. That is, the algorithm fails to compute the correct neck-shaft angle for only 0.7% of the samples.

5 Conclusion

A method of computing neck-shaft angle for detecting femur fracture is presented in this paper. Given the contour of a femur, level lines that are perpendicular to the contour are computed. The mid-points of the level lines at the shaft are aligned with the shaft axis. The largest bundle of long level lines at the head gives an approximation of the neck axis. Given this approximation, an optimzation algorithm is applied to find the best-fitting axis of mirror reflection of the head-neck contour. This axis of mirror reflection is the best-fitting neck axis. The neck-shaft angle can now be computed from the neck and shaft axes. Test results show that the algorithm correctly computed the neck-shaft angles for 99.3% of

Table 2. Summary of misclassifications.

wrongly classified as healthy		
type	training	testing
change in shape	2	6
no change in shape	2	12

wrongly classified as fractured		
type	training	testing
no shaft	0	3
unusual pose	2	1
program failure	1	2

the training and testing images. Application of the computed neck-shaft angle for fracture detection achieved an accuracy of 94.4% for training images and 92.5% for testing images. We are investigating other fracture detection methods to complement the current method and to improve the detection accuracy.

Acknowledgment

This research is supported by NMRC/0482/2000.

References

1. IOF: The facts about osteoporosis and its impact. International Osteoporosis Foundation, http://www.osteofound.org/ press_centre/fact_sheet.html (2002)
2. Ryder, D.M., King, S.L., Olliff, C.J., Davies, E.: A possible method of monitoring bone fracture and bone characteristics using a non-invasive acoustic technique. In: Proc. Int. Conf. on Acoustic Sensing and Imaging. (1993) 159–163
3. Kaufman, J.J., Chiabrera, A., Hatem, M., Hakim, N.Z., Figueiredo, M., Nasser, P., Lattuga, S., Pilla, A.A., Siffert, R.S.: A neural network approach for bone fracture healing assessment. IEEE Engineering in Medicine and Biology **9** (1990) 23–30
4. Singh, V.R., Chauhan, S.K.: Early detection of fracture healing of a long bone for better mass health care. In: Proc. Annual Int. Conf. of IEEE Engineering in Medicine and Biology Society. (1998) 2911–2912
5. Laugier, P., Padilla, F., Camus, E., Chaffai, S., Chappard, C., Peyrin, F., Talmant, M., Berger, G.: Quantitative ultrasound for bone status assessment. In: Proc. IEEE Ultrasonics Symposium. (2000) 1341–1350
6. Matani, A., Okamoto, T., Chihara, K.: Evaluation of a trabecular structure using a multiresolution analysis. In: Proc. Annual Int. Conf. of IEEE Engineering in Medicine and Biology Society. (1998) 632–633
7. Tascini, G., Zingaretti, P.: Automatic quantitative analysis of lumbar bone radiographs. In: Proc. IEEE Nuclear Science Symposium and Medical Imaging Conference. (1993) 1722–1726
8. Furst, J.D., Pizer, S.M.: Marching optimal-parameter ridges: An algorithm to extract shape loci in 3D images. In: Proc. Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (LNCS 1496). (1998) 780–787
9. Guo, D., Richardson, P.: Automatic vessel extraction from angiogram images. In: IEEE Conf. on Computers in Cardiology. (1998) 441–444
10. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. Int. Journal of Computer Vision **1** (1987) 321–331
11. Xu, C., Prince, J.L.: Gradient vector flow: A new external force for snakes. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition. (1997)
12. Leow, W.K., Li, R.: Adaptive binning and dissimilarity measure for image retrieval and classification. In: Proc. IEEE CVPR. (2001)