

Segmentation of Cam-Type Femurs from CT Scans

Gabriel Telles O'Neill, Won-Sook Lee

Abstract

We introduce a new way to accurately segment the 3D femur from pelvic CT scans. The femur is a difficult target for segmentation due to its proximity to the acetabulum, irregular shape and the varying thickness of its hardened outer shell. We overcome these difficulties by (a) dividing the femur into two rounds of segmentation - one for the femur head and another for the body - (b) pre-processing the CT scan to reduce anatomical sources of error (c) two modes of segmentation - a rough estimation of a contour and another for fine contours. Segmentations of the CT volume are performed iteratively, on a slice-by-slice basis and contours are extracted using the morphological snake algorithm. Our methodology was designed to require little initialization from the user and to deftly handle the large variation in femur shapes, most notably from deformations attributed to cam femoral-acetabular impingements. Our efforts are to provide physicians with a new tool that creates patient-specific and high-quality 3D femur models while requiring much less time and effort. Femur models segmented with our method had an average volume overlap error of $2.71 \pm 0.44\%$ and symmetric surface distance of $0.28 \pm 0.04\text{mm}$ compared to ground truth models.

Keywords: Segmentation; Femur; Femoral-Acetabular Impingements; Hip; Morphological Snake

1 Introduction

In medical imaging, segmentation can be used to extract contours of organs or bones, which can later be used by physicians to aid in their diagnosis. Femoral-Acetabular Impingements (FAIs) refer to a source of hip pain where deformities exist on the hip bones which deteriorate the joint's soft tissue over time. Segmenting a pelvic Computed Tomographic (CT) scan allows for the creation of patient-specific hip-bone models. These models can then be used by physicians to detect the locations and severity of impingements as well as to aid in the planning of minimally invasive procedures to correct the condition.

We are motivated by the medical need for patient-specific 3D models of hip bones for the treatment of FAIs. At present, the most precise results for pelvic bone segmentations from CT scans come from the intensive manual labeling of voxels. While providing the most accurate results, manual segmentation is extremely time-consuming; requiring hours of work for a single pelvic scan, even with computer assisted functions. Our goal is to create a machine-driven segmentation scheme which greatly reduces the time required to segment the femur from a pelvic CT scan while returning the sufficiently accurate results required for physicians to treat FAIs.

There are number of factors which complicate attempts to segment a femur from a CT pelvic scan. These complications can be classified as either technical limitations associated with CT imagery (such as low resolution, artifacts and lack of color information), patient-specific complications (such as bone size, varying caput-collum-disphyseal angles and the potential presence of FAI), or complications particular to segmenting the femur (such as its close proximity to the

acetabulum, it's off-tubular shape or the inconsistent density of its osseous tissue).

Our contribution consists of providing a segmentation procedure which overcomes or mitigates these complications. The human hip has bones whose shapes and size can vary greatly between patients. Regional shape differences are especially expected for patient's suffering from hip illnesses, like FAIs. We resolve to create a practical solution to femur segmentation which works across a spectrum of shapes and validate our method on a diverse sampling of patient scans. Our solution subdivides the femur into two, partially overlapping regions for segmentation: the femur-head and the femur-body. For both of these objects, we extract the contours from cross-sections of the femur as they appear on singular slices from the CT volume. While we use Morphological Snakes (Álvarez, Baumela, Henríquez, & Márquez-Neila, 2010) as our contour extraction method due to its stability and improved speed, our main contributions lies in the methods we use to subdivide femur segmentation into smaller tasks and the pre-processing steps we apply prior to segmentation. Both the subdivision and the preprocessing methods stem from our analysis of inter-patient variability in terms of bone structure and density.

We pre-process the contents of the CT pelvic scan by (a) cutting-out portions of the acetabulum & filling its space with a low-gradient fill, (b) upscaling the voxel's resolution, (c) performing ceiling thresholding to emphasize voxels composed of compact bone and (d) applying a flooring thresholding to flatten the gradients between of soft tissues.

The structure of this paper is a follows: Section 2 contains background information, Section 3 describes our femur segmentation strategy, Section 4 contains some results and discussion and Section 5 contains closing remarks and directions for future work.

2 Background

2.1 The Hip Joint

The *acetabulofemoral joint* exists between the pelvis and the thigh and is composed of two bones; the *femur* and the *acetabulum*. Between these bones are soft tissues, such as the acetabular labrum and articular cartilage, which protect the contact surfaces of the bones and facilitate the smooth, pain-free operation of the joint



Figure 2.1 Pelvic x-ray with highlighted femur (red) and acetabulum (blue)

The upper extremity of the femur has a number of notable features which affect segmentation:

- The *femur head* is the rounded end of the femur. Its surface comes in the closest contact with the pelvic bones and can generally be approximated with the shape of a sphere or conchoid
- The *pit* (or fovea) for the ligament of the head of the femur is an indentation in the otherwise spherical surface of the femur-head
- The *femur body* is the tubular shaft which makes up the majority of the femur
- The *greater trochanter* is a large bony-bump at the top of the body, opposite to the femur head

CT images of the femur-head's subchondral bone composition show large density variations, like a very thin compact shell compared to the much thicker one found in the femur body. Comparatively, regions of the acetabular rim covering the femur head are also noticeably denser.

The hip-joint's task in the human body is to support upper-body weight- both while standing still and during movement. Hip movement is enabled by the hip functioning as a ball-and-socket (Field & Hutchinson, 2008) where the femur head constitutes the *ball* and the acetabulum provides the retaining *socket*.

2.2 Femoral Acetabular Impingement

Hip impingement, or *Femoral-Acetabular Impingement* (FAI), is a pathological condition where there is a deformity on either of the hip-joint's bones. This deformity usually manifests itself as a bony-bump on the acetabular rim (*Pincer impingements*) or the femur's head-neck junction (*Cam impingements*) or both. As a result, the hip joint will lose its ideal ball-and-socket shape which causes abnormal contact (impingement) between the femur and acetabulum during normal hip-rotation. In turn, this abnormal contact causes chaffing of the soft-tissue protecting the two bones, deteriorating the integrity of the joint over time. FAI is most often associated with pain during hip flexion, adduction, and the internal rotation of the femur. If left untreated, a FAI can lead to cartilage damage, labral tears, early hip arthritis, hyperlaxity, sports hernias, chronic lower back pain (Hossain & Andrew, 2008) and an eventual hip replacement surgery.

2.3 Medical Segmentation Methods

Medical images are a popular proving ground for segmentation methods as the images are usually full of complex shapes, noise and the extracted shapes have obvious, life-changing uses.

The original energy minimizing curves, or "snakes" (Kass, Witkin, & Terzopoulos, 1987), are perhaps the best known segmentation scheme. Since their introduction in 1987, a number of modifications/enhancements have been proposed. These include the addition of a balloon force operator (Cohen L., 1991), 3D generalizations of the snake model (Cohen & Cohen, 1993), likening snakes to level-set methods with geodesic active-contours (Caselles, Kimmel, & Sapiro, 1997), increasing the a snake's capture range & ability to evolve into concavities (Xu & Prince, 1998) and add topological flexibility (McInerney & Terzopoulos, 2000). Model-based segmentation schemes are useful in cases where part of an object's information is missing. Early cases include Active shape models (Cootes, Taylor, Cooper, & Graham, 1995) and active appearance models (Cootes, Edwards, & Taylor, Active Appearance Models, 1998) based around point distribution models.

Similar to energy-minimizing curves are methods which are region-based rather than boundary-based. A prime example of this is the Chan-Vese method (Chan & Vese, 2001) (Vese & Chan, 2002) which seeks to minimize the energy inside a curve through the Mumford-Shaw functional. At roughly the same time, Diffusion Snakes (Cremers, Schnörr, Weickert, & Schellewald, 2000) were detailed, which used prior shapes along with the Mumford Shaw functional. As an aside, a model-based method which originally used the kernel density information of shape-priors (Cremers, Osher, & Soatto, Kernel Density Estimation and Intrinsic Alignment for Shape Priors in Level Set Segmentation, 2006) was modified to also include intensity-priors or regions inside the boundary (Chen & Radke, 2009).

A number of segmentation strategies have been proposed focusing on the segmentation of femurs from MRI and CT scans, although none of these address the much more difficult task of segmenting hips suffering from FAIs. One such method required assigning a scan into one of four groups depending on the anticipated difficulty of segmentation, and in the worst case, separating the femur and acetabulum using a combination of the Hueckel operator and orthogonal line detection (Zoroofi, et al., 2003). Unfortunately, their techniques returned many moderate and poor results. Another method involved a significant amount of user-input, requiring the user to manually surround the femur head with contour points for a snake method and making manual correction whenever the snake failed due to soft edges (Magenat-Thalmann, Yahia-Cherif, Gilles, & Molet, 2003).

In regards to model-based segmentation methods for femur bones, coarse-to-fine methods have been used with 3D meshes (Gilles, Moccozet, & Magneat-Thalmann, 2006) for anatomical modeling, MRI scans with low resolutions or fields of view (Schmid, Kim, & Magneat-Thalman, 2011), and statistical shape models (Yokota, Okada, Takao, Sugano, Tada, & Sato, 2009) have been used to segment diseased hips. In addition, 2D point-distribution models for ASMs (Song, Li, Ou, Han, Zhao, & Wang, 2007) have been tested against healthy hips. Generalized models of hip-bones can lead to complications when the object of segmentation has a FAI. This is due to the bony bump being outside the model's expected distribution. Conversely, models specifically tailored for bones with impingements require a very large set of prior shapes as the location of the bony bumps can be highly irregular. To return the best segmentation results, a method which is elastic to the pronounced differences between bones is required.

2.4 Morphological Snakes

Our segmentation solution selected a 2D implementation of Morphological Snakes (Álvarez, Baumela, Henríquez, & Márquez-Neila, 2010) to extract our desired contours. This method is a recent modification of the well-known Geodesic Active Contours. The major difference between the two models is how each method solves the partial differential equations (PDEs) responsible for curve evolution. While Geodesic Active Contours expresses these PDEs with a set of differential operators, Morphological Snakes instead takes the approach of approximating these terms as the composition of morphological operators. Specifically, the inf (infimum) and sup (supremum) operators. This substitution claims three advantages:

1. **Simplicity of Implementation** – The level-set is expressed as a binary piecewise constant function

2. **Speed** – Execution time in test was shorter, sometimes by full order of magnitude
3. **Stability** – No re-initialization of the level set or contour is required

3 Segmentation of the Femur

3.1 User Initialization

Our segmentation strategy is considered semi-automatic instead of fully automatic because it requires preliminary user-interaction in order to achieve its goal. The current stage of the segmentation process contains all of these interactions, allowing the following stages to perform automatically.

The user’s input is required in order to demarcate 3D regions-of-interests (ROIs) within the volume of the CT scan. We segment the femur head and body semi-independently, which requires two partially-overlapping ROIs to be specified before the segmentation procedure can begin in earnest.




Figure 3.1 Frontal view of the ROI for the femur (left) head (right) body with split slice in yellow

The first of these ROIs encompasses the whole femur head, the upper portion of the greater trochanter and neighboring sections of the acetabulum. We refer to it as ROI_H . Its vertical upper-bound is the topmost slice containing traces of the femur head while its lower-bound contains the bottom-most traces of the femur head’s sphericity as shown in Figure 3.1. In addition to this, a voxel whose location roughly constitutes the center of the femur head is also attached to this ROI.

The second ROI encloses the more tubular shape of the femur body up to the top of the greater trochanter as shown in Figure 3.1. We refer to it as ROI_B . Attached to ROI_B is the index of the slice where the cross-sections of the femur first appear as two disjointed objects (the greater trochanter and femur head), which we call Z_S .

3.2 Acetabular Rim Removal

The primary driving force behind our segmentation strategy is the morphological snake algorithm. If a slice contains a section of the femur, we can envelop it with an initial contour for the snake method. The snake method will then be attracted to the femur’s edges and provide us with an accurate outline of the femur whose volume we can extract. However, in cases where the femur is near to another object with strong edges, the snake may mistakenly be attracted to its edges instead, which leads to much less accurate outlines of the femur head. This is the case for regions of the femur head being cupped by the acetabulum.

To neutralize this unwanted source of competing edges, an “acetabular-removal” step is added to our solution, which is made easier by the fact that the acetabular rim has much

denser bone than can be found on the femur head (leading to better defined edges which aid in segmentation). We cut-out portions of the acetabulum from our source data and replace the resulting hole with a smooth, muscle-like fill. Our starting point for this procedure is ROI_H .

Thresholding

Our first step towards removing the acetabulum is to perform a thresholding of the image to isolate segments of compact-bone. Sections of the acetabulum adjacent to the femur have a shell of compact-bone thick enough to preclude weak edges. As a result, thresholding a slice with an appropriate value provides a solid outline of the acetabulum in one or two connected fragments.

Creating the Mask

After having retrieved a binary-image mapping the location of compact bone voxels, we create a mask designed to cover the acetabulum only. To do this, we decompose the thresholded image into sets of connected white voxels and filter-out fragments which do not meet the following criteria:

- a) A fragment must be composed of a sufficient number of white voxels
- b) A fragment’s location must be on the expected side (right or left) of the femur head center
- c) A fragment must have all of its elements a sufficient distance away from the femur head center

Any fragments which weren’t filtered-out get imprinted onto a mask for the acetabulum. Following this, a dilation operation is applied to the mask with a 3×3 square operator. This has the advantage of both joining disconnected fragments and also capturing the softer outline of the acetabulum which wasn’t captured during the thresholding step.

The final touches include flood-filling a tertiary color into the corners of the mask which allows us to distinguish between the empty area around the mask and the spongy-bone regions inside the acetabulum.

Filling the Mask Region

The final step of the acetabular removal process consists of replacing voxels which the mask identifies as belonging to the acetabulum. From the start, the goal of the acetabular removal was to eliminate some of the gradient which might cause the snake segmentation method to return erroneous results.




Figure 3.2 (a) Original Slice of the femur-head (b) Slice of the femur-head after thresholding (c) Mask identified for acetabulum (d) Identified mask after dilation & flood-filling (e) Mask-region with black fill (f) Mask region with low-gradient fill

To this end, we have opted to not replace the acetabulum with a flat color, but with a low-gradient fill instead. This low-gradient fill was designed to provide a smooth appearance at all points of the cut-out acetabulum. It works as an offshoot of Gaussian smoothing, repurposed to fill in the acetabular cut-out. Low gradient fill starts at the outside

of the cut-out and works its way in. An additional feature in the algorithm is that it converge the fill to a muscle-like intensity in order to prevent the formation of high-gradient spines along the center of the cut-out. The entire process can be seen in Figure 3.2.

It is important to note that acetabular removal is not concerned with segmenting the whole acetabulum. It instead concentrates on identifying parts of acetabulum which are near to the femur head and marginalizing their impact over the snake-based segmentation methods.

3.3 Pre-processing

In addition to acetabular rim removal, we use three pre-processing methods which generally improve the quality of our results: upscaling, ceiling thresholding and flooring thresholding. By default, all three methods are applied prior to each one of the contour-extraction steps. However, these functions can also be toggled-off individually and for individual steps of the segmentation procedure if a user has reason to believe that it will improve results.

3.3.1 Upscaling

The internal forces acting on a snake’s evolving contour exist to ensure that the curve remains smooth and well-rounded. This is desirable for segmenting objects whose curvature arc slowly. However, when applied to objects with narrow sections or sharp changes in the direction of edges, this internal force can be the source of premature termination of curve evolution or cause the curve to enter the object it is attempting to segment (He, et al., 2008).

Due to the comparatively low resolution of the CT slices, small concavities exist in the in cross-sections of the femur head and neck which may invoke the detrimental effects of an evolving snake’s internal forces.

A solution which has improved results is to *upscale* the contents of the ROI. In this case, upscaling refers to resizing a slice’s contents inside the ROI by an integer scale factor, without using any interpolation. This has the effect of artificially widening narrow concavities on the femur’s surface and diluting sharp angles in a cross-sections contour. Once a final contour has been extracted from an upscaled ROI, the contour is downscaled by the same scale factor prior to becoming our final segmentation result.

The downside to this solution is that it increases the amount of time required to extract a femur contour. This can be attributed to the size of the evolving contour being increased, resulting in needing to calculate the external forces being exerted on the curve by more voxels. However, we find this trade-off acceptable for small scale-factor values due to the great improvement in results.

3.3.2 Ceiling Threshold

We apply a *ceiling threshold* method to highlight bone voxels while still keeping lower-intensity voxel information as shown in Figure 3.3. The ceiling threshold value is chosen to be sufficiently high as to only highlight voxels which are guaranteed to correspond to compact bones. The highlighted voxels are nearly guaranteed to stop the snake’s evolving contour.




Figure 3.3 Slice (left) before ceiling-thresholding (right) after ceiling-thresholding

3.3.3 Flooring Threshold

We can apply a *flooring threshold* method to unify the intensities of soft tissues, as seen in Figure 3.4. Primarily, this method is used to remove local minima caused by streaks of fat between muscles in the thigh. It should also be noted that for any slice of ROI_H which has a flooring threshold applied, the acetabular rim removal step can be modified to omit the low-gradient fill phase. This fill would be floored, which negates its usefulness.




Figure 3.4 Slice (left) before flooring-thresholding (right) after flooring-thresholding

3.4 Segmenting the Femur Head ROI

The femur head is near-spherical in shape. Correspondingly, a cross-section of the femur-head which appears in CT slices will appear near-circular. During the evolution of a snake’s contour, internal forces assure that the curvature of the contour remains smooth. A side-effect of these internal forces is that snake methods cannot capture contours which have sharp corners or narrow protrusions. The near-circular cross-section of a femur head, ideally possessing neither of these features, is therefore a good candidate for snake-based segmentation methods.

Our segmentation strategy for the femur head consists of two consecutive steps: (a) a “primer-slice” is segmented, which is the first slice to be processed in the ROI and then (b) the subsequent slices are segmented, a process which starts at the primer slice’s neighbours, and cascades outwards.

3.4.1 Primer Slice Segmentation

We refer to the slice in an ROI which is the first to have its femoral contour extracted as the “primer” slice. What differentiates this slice’s segmentation process from the others is how its initial contour is derived. Unlike the procedure for other slices, the primer slice has its femur head contour segmented twice: the first time to attain a *rough* contour and the second time to achieve a semi-final *fine* contour. This twice-over segmentation is a cost-effective way to free the user from setting an initial contour.

The *rough* contour extraction starts with an initial contour the size of the whole ROI. When used as input for the first morphological snake pass, this oversized initial contour will deflate into an approximation of the femur head or body which may have a few defects. Afterwards the *fine* contour extraction commences; the final contour obtained through the rough segmentation is re-processed to be used as the initial contour for the fine segmentation.

For the ROI_H , the primer slice is chosen to be the slice which is nearest to the center. Formally, the primer slice Z_p is chosen with Eq. 3.1, where Z_t is the top-most slice of ROI_H and Z_b is its bottom-most slice:

$$Z_p = Z_t + \text{round}\left(\frac{Z_b - Z_t}{2}\right) \quad \text{Eq. 3.1}$$

Finding the rough contour

The rough contour extraction serves to find a starting area for the fine contour extraction. The complete process of obtaining the final rough contour can be seen in Figure 3.5. The initial contour for this stage is the largest circle whose entirety fits inside ROI_H 's area. This circle has its center at the ROI's mid-width and mid-height. The circle will also have a radius equal to the ROI's mid-width or mid-height; whichever is smallest. The femur-head is bound to be fully within the limits of the circle due to the landmarks used to size the ROI_H .




Figure 3.5 Primer slice for rough segmentation (a) as it originally appears (b) after acetabular removal (c) with super-imposed initial contour (d) with superimposed final contour (e) with super imposed circle from Hough transformation (f) with super-imposed widened circle

This circular initial contour is then evolved using the morphological snake algorithm (Álvarez, Baumela, Henríquez, & Márquez-Neila, 2010). For this first pass, the snake parameters are set to have a strong, negative balloon force and a weak edge-attraction force.

Given the circular initial contour and the right parameters, the morphological snake will return a rough final contour which approximates the femur head. This contour may approximate the femur-head's exterior well or it may contain

a few flaws. For the rough contour, a few protrusions or depressions from the femur head's shell is acceptable. What's most important to retrieve from the first pass of the morphological snake algorithm is the general curvature of the femur head.

The final step in the rough contour extraction is to provide an initial contour for the fine contour extraction. We have chosen to use a close-fitting circle to this end. To get a circle from our rough final contour we:

1. Perform a Hough circle transform on the rough final contour to determine the circle which best approximates the shape of the femur. We improve our results by using an elliptical-Gaussian kernel voting scheme similar to the one used by (Fernandes & Oliveira, 2008). In addition, we use a narrow-band approach around the ROIs femur-center point to speed-up the process.
2. Widen the circle's radius in order to ensure all of the voxels belonging to the femur-head are included inside the initial contour

The contour re-initialization has proven useful in scenarios where distant fragments of the acetabulum linger after the rim is removed or the greater trochanter's top appears in the primer slice. These scenarios can cause small sections of the contour to remain isolated from the femur head during curve evolution. Re-initializing the contour through Hough circle detection allows us to guarantee that all sections of the initial contour for the fine step are within the same range of the femur head.

Finding the fine contour

The fine contour segmentation process constitutes the second-pass of the morphological snake on the primer slice. During this pass, the snake operation provides us with the decisive segmentation result for this particular slice. The detected circle with widened radius is used as our initial contour.

The snake parameters used in the second pass focus on a strong attraction force with a weaker deflation force. This is contrary to the focus of the first snake's pass. This change is validated by our knowledge that the initial contour will be much closer to the desired final contour in the second pass than the first.

Because of these changes, the snake algorithm will return results with are tighter to the bone and less-error prone than those of the rough contour. An example of the final segmentation of the fine contour can be seen in Figure 3.6




Figure 3.6 Example final segmentations for the femur head detected during the (left) rough and (right) fine contour extraction

3.4.2 Subsequent slice segmentation

Once the primer slice has been segmented, we can begin sequentially segmenting the other slices in ROI_H . The slices are queued for segmentation based on their distance from the primer slice. If the primer slice is noted as Z_p , then the next slices to be segmented will be Z_{p+1} & Z_{p-1} , followed by Z_{p+2} & Z_{p-2} , Z_{p+3} & Z_{p-3} and so on. Segmenting the whole ROI

entails performing individual, sequential segmentations away from the primer slice in both the superior and inferior direction, ending when the top-most and bottom-most slices in the ROI have been segmented.

The method of segmentation for these slices is similar to the one reported during the fine contour stage. They differ only in how they obtain their initial contour: we have decided to base the initial contour of each slice on the final contour of its neighboring slice. Thus, an un-segmented slice Z_i will base its initial contour on Z_{i+1} if it is above the primer slice and Z_{i-1} if it is below the primer slice. However, before a neighbor's final contour can be used as an initial contour, it is expanded. This expansion is used to safeguard against moderate differences in shape between neighboring femur-head cross sections. The flow can be viewed in Figure 3.7 .




Figure 3.7 (a) Slice Z_i in need of an initial contour (b) Slice Z_{i+1} with final contour (c) Slice Z_{i+1} with dilated final contour (d) Slice Z_i using Z_{i+1} 's dilated final contour as an initial contour

All that remains to obtaining a final contour is to supply the initial contour to the morphological snake algorithm and evolve it using the "fine" contour extraction parameters.

3.5 Segmenting the Femur Body ROI

The femur, like most long bones, has as much thicker shell of compact osseous tissue around its shaft than it does at its epiphyses (Rizzo, 2001). This results in the slices from the femur body being easier to segment than those from the femur head, on account of the compact bone's high contrast with the soft tissue that surrounds it. In addition, the femur body does not have adjacent bones which need to be isolated due to their proximity, once again unlike the femur head with the acetabulum.

The femur body, being vertically tubular in shape is another good candidate for a slice-by-slice segmentation strategy as was used for the femur head. In fact, the contour extraction method we use for the femur body follows closely the one employed for the femur head, with a few minor exceptions. The biggest of these exceptions being that ROI_B has Z_s ; a slice where the femur splits into two disjointed objects. This split will necessitate a single-slice exception on how initial contours are derived.

Our segmentation strategy for the femur body once again consists of two consecutive steps: (a) a primer slice has its rough and fine contour extracted (b) the segmentation of

subsequent slices, where the primer slice's segmentation initiates an outward segmentation cascade.

3.5.1 Primer Slice Segmentation

For the ROI_B, the primer slice is chosen as the bottom-most slice, which we refer to as Z_b . This slice should also correspond to the bottom-most slice of the whole pelvic scan, and contain a tubular cross-section of the femur body.

Finding the rough contour

Once again, the rough contour extraction finds a starting area for the fine contour extraction. The initial contour for this stage is the largest circle whose entirety fits inside the ROI_B's area. Once this has been fitted, the morphological snake algorithm evolves the contours using the same set of "rough" parameters used on the femur-head's primer slice.

Given the thick, high-contrast shell of the femur body, the final contour should be virtually free of defects and quite close to the result that will be obtained during the fine contour extraction phase (see Figure 3.8). Slightly jagged contours are likely to be the only visible flaws on the rough contour, and even these will be reset using the fine step.




Figure 3.8 (left) Initial contour for rough step (right) final contour for rough step

Finding the fine contour

The fine contour serves the same purpose here that it did for the femur head: providing the decisive segmentation of the primer slice. Unlike the femur-head however, the cross-section of the femur body is not near-circular, meaning we do not employ a Hough transformation on the rough step's final contour to extract the best-fitting circle. Instead, we expand the rough step's final contour much as we do for the sequential slices. The result is a curve which is near evenly-distanced from the femur body at the primer slice (see Figure 3.9). We use this contour as the fine step's initial contour

The final contour for the primer slice is obtained when the morphological snake algorithm is applied to the new initial contour. Correspondingly, we use the "fine" set of segmentation parameters to ensure the best results.




Figure 3.9 (left) Initial contour for the fine step (right) Final contour for the fine step

3.5.2 Subsequent Slice Segmentation

With final contour of the primer slice well defined, we continue with the segmentation of all the other slices in

ROI_B. Our procedure for doing so is familiar: performing successive segmentations away from the primer slice. However, the primer slice for the ROI_B is the bottom-most slice, so we can only iterate upwards.

Thus, an un-segmented slice Z_i will always base its initial contour on Z_{i+1} . However, prior to using Z_{i+1} 's final contour as Z_i 's initial contour, we expand the contour, ensuring that no sections of the femur body are outside its domain.

However, there exists a single slice in ROI_B which requires exceptional initialization. That slice is Z_s ; where the greater trochanter and femur-head first become two, disjointed objects. The split occurs at the height of intertrochanteric line and the greater trochanter crowns shortly thereafter. Of the two objects, we focus on segmenting the greater trochanter as the femur head has already been segmented from ROI_H.

The initial contour for Z_s is based not only the final segmentation of slice Z_{s+1} in ROI_B, but also from the final segmentation of slice Z_{s+1} in ROI_H. For shorthand's sake, we refer to the set of voxels contained within the final contours of slice Z_{s+1} for ROI_H as V_H and the set of voxels contained within the final contours of slice Z_{s+1} for ROI_B as V_B . By subtracting the set of voxels in V_H from V_B , we can obtain a good estimation of the greater trochanter's cross-section in slice Z_s . Thus we use the largest set of connected voxels of $V_B - V_H$ to provide an initial contour for slice Z_s . (See Figure 3.10).




Figure 3.10 (a) Final contour for slice Z_{s+1} in ROI_B (b) Voxels of set V_B (c) Final contour for slice Z_{s+1} in ROI_H (d) Voxels of set V_H (e) Volume for ROI_H subtracted from ROI_B (f) slice Z_s with its initial contour

No matter how one of the subsequent slices obtains its initial contour, it is used as input for the morphological snake algorithm along with the set of “fine” snake parameters. The final curve provided by the algorithm becomes our definite segmentation for ROI_B.

3.6 Post-Processing

The final step in our segmentation process is the creation of a single model for the segmented femur. This is resolved by calculating the union between the voxels encased by ROI_H's final contours with voxel's inside for ROI_B's final contours. This union will allow us to consolidate the slices which are only covered by one of the ROIs versus those covered by both as seen in Figure 3.11. From here, the results can be viewed as its native collection of voxels or be polygonized for an improved presentation.




Figure 3.11 Sample union of the 3D model from ROI_H and the 3D model from ROI_B

4 Results and Discussion

In order to assess the quality of our femur segmentation strategy, we compare the volume segmented by our semi-automatic method to the volume of a manual segmentation. The manual segmentation's volume was labeled on in a slice-by-slice, voxel-by-voxel basis using an image editing program. It can be considered the “ground truth” segmentation. Our goal in testing is to obtain semi-automatic results which are identical, or within human error of this ground truth.

4.1 Hospital Database

We have collected 20 pelvic scans from a study into FAIs conducted by the Ottawa General Hospital. Of these scans, 17 were from CAM type patients (7 bilateral and 10 unilateral) and 3 were part of the control group. For each patient we also have one antero-posterior and one lateral radiograph. In conjunction, the scans and radiograph have allowed us to study the variations between patients.

4.2 Testing Database

From these 20 scans we selected 6 that provided as diverse a sampling as possible and manually segmented both of their femurs for validation purposes. The clinical characteristics of these patients can be seen in Table 4.1.

Table 4.1 Clinical characteristics of the six patients whose femurs were manually segmented

Patient ID #	Age	Sex	FAI Diagnosis
1	39	M	Bilateral
2	47	F	Left
3	29	F	Left
4	33	M	Left
5	22	M	Right
6	54	M	Control

Other than femur size, orientation and the presence of FAIs, our sample set shows a number of other inter-patient variations which would test the rigour of any femur segmentation strategy. The most important dissimilarities

between patients are inevitably located at or near the femur head. The femur body's thick, tubular shape lessens the impact of inter-patient differences as far as segmentation is concerned.




Figure 4.1 Femur and acetabulum appear (left) well-spaced (right) close at points

As an example, Figure 4.1 shows a slice where the boundaries of the femur-head and acetabulum appear quite close due to blurred edges and CT noise. Any segmentation strategy will need to separate these two objects.




Figure 4.2 Cross-sections where the femur head appears (left) near-circular (right) deformed

Strict model-based methods will have difficulty getting the contours from the right image in Figure 4.2 as it's a far cry from the average femur shape. Any segmentation strategy which seeks to enforce a sphericity constraint over the femur head will especially have difficulty here.




Figure 4.3 Concavity between femur head and greater-trochanter is (left) wide (right) narrow




Figure 4.4 Various shapes of the femur head's fovea

Differential gaps between femur head and greater trochanter can be seen in Figure 4.3, where the more narrow of the two is expected cause some snake methods to halt their curve evolution prematurely. Also worth noting is the higher prevalence of compact bone in the slice with the narrow concavity.

Finally, Figure 4.4 illustrates a wide array of foveas. Some of which are nearly straight, other are angular concavities while others still appear as deep dimples.

4.3 Testing Conditions

All of our scans were acquired using a GE Lightspeed Plus. For each one of them, the slices have a resolution of 512 by 512 voxels, where each voxel had a physical dimension of 0.68 x 0.68 x 1.25 mm.

As for parameters used during the preprocessing step, we started off with an *upscale factor* of 2. The resolutions of the resized-slices were sufficiently large that the evolving snake was better able to map the concavities on the interior of the greater trochanter. In addition, the upscale factor gave a nicer edge to some of the sharp turns in the femur head's exterior.

During the acetabular removal step, we used a default threshold value of 316 HU (for all tested femurs) to highlight voxels which belonged to the acetabulum's hard shell. Across all six CT scans, we found that this value did well to emphasize the acetabulum's bones without accidentally accentuating the voxels for soft-tissue between the acetabulum and femur head.

We also applied a flooring threshold exclusively to the slices in ROI_B with default value of 127 HU in order to flatten the intensities of muscle and fat voxels near the bone.

We used a default value of 201 HU (once again, for all the tested femurs) during the ceiling threshold step to highlight voxels corresponding to compact osseous tissue. After some experimentation, we discovered that lowering the ceiling threshold for the top-most slices of the femur head would occasionally improve results. For this reason we use a threshold value of 144 HU for the top-most 25% of slices above the primer slice. This accounts for the top-most 12.5% of the whole femur head ROI.

The *rough* and the *fine* contour extractions methods use two different sets of procedures, as well as input parameters. The *rough* method uses an edge-based balloon stopping criterion whereas the *fine* method uses an image intensity-based balloon stopping criterion. The parameters for the former can be found in Table 4.2 and the latter in Table 4.3.

Table 4.2 Test parameters for Morphological snake with edge-threshold-balloon (Rough Contours)

σ	N	ETB	EDT
0.0	100	0.1	1.000

Table 4.3 Test parameters for morphological snake with snake-balloon-difference-radius (Fine Contours)

σ	N	DR	EDT
0.0	100	12	0.001

For these parameters, $G(I)$ is the level-set method the contours are trying to minimize, σ is the standard deviation of the Gaussian convolution applied to the original image, N is the maximum number of iterations, EDT is the threshold for $G(I)$ during the speed propagation procedure, ETB is the threshold for $G(I)$ during the balloon propagation procedure, and DR is the difference magnitude with respect to the central pixel value that propagates the balloon.

4.4 Qualitative Segmentation Comparison

The images found in Table 4.4 and Table 4.5 originate from our patient #3's left hip, which was diagnosed with a

cam impingement. The CT scan on which the images are based is composed of 126 slices which spanned from the top of the patient’s sacrum to the roughly a 1cm below her minor trochanters.

Table 4.4 Comparison of our segmentation method results for the femur head with those from a manual segmentation

	Our Method	Ground Truth	Difference
Head Top			
Head Mid			
Head Bottom			

A side-by-side comparison of segmentation results obtained through our method and those obtained through manual segmentation can be seen in Table 4.4 for the femur head and Table 4.5 for the femur body.

Table 4.5 Comparison of our segmentation method results for the femur body with those from a manual segmentation

	Our Method	Ground Truth	Difference
Greater Troch.			
Minor Troch			
Femur Body			

The left column contains the final contour of our method for a particular slice. Volume highlighted in white represents voxels which inside the final contour evolved by morphological snakes.

The middle column contains the volume which is considered the ground truth segmentation. Volume highlighted in white was manually segmented on a slice-by-slice, voxel-by-voxel basis with an image editing program.

The right column shows how well these two segmentations overlap. Voxels shown in *black* indicate identical results for both segmentations. Green voxels indicated voxels which were a member of the ground truth segmentation but not of our segmentation while *red* voxels indicate the opposite.

Table 4.6 Volumetric renderings of the first three patients’ left femurs

ID #	Our Method	Ground Truth
1		
2		
3		

Table 4.7 Volumetric renderings of the last three patients’ femurs

ID #	Our Method	Ground Truth
4		
5		
6		

In order to provide a more complete visual assessment of how the models obtained through our method compare to the ground truth models, 3D volume renderings have been

obtained for each model. The left femurs from the first three patients can be found in Table 4.6 while the right femurs from the last three patients are in Table 4.7. The middle column of each table contains the renderings of femurs segmented with our method while the last row contains renderings from manually segmented femurs. All renderings were taken from an antero-posterior (AP) position.

The technique (Stegmaier, Strengert, Klein, & Ertl, 2005) used to produce these 2D projections is an example of a Digitally Reconstructed Radiograph (DRR) (Sherouse, Novins, & Chaney, 1990). DRRs are an approximation of a 2D x-ray image obtained from CT or MRI data. A DRR's semi-transparent approach to volume rendering allowed us to present more depth compared to polygonized surface models.

4.5 Quantitative Segmentation Comparison

The accuracy of our segmentation method is measured numerically in terms of the intersection of voxels covered by both models and the distance from the bordering voxels in the ground truth model to the nearest bordering voxels in the model segmented through our method. The volumetric overlap errors (VOEs) of the models can be seen Table 4.8. This table displays the global error (VOE_g), the average slice error ($VOE_s Avg$), the slice minimum error ($VOE_s Min$) and slice maximum error ($VOE_s Max$).

Table 4.8 Volumetric overlap error

ID #		VOE_g (%)	$VOE_s Avg.$ $\pm stdev$ (%)	VOE_s Min (%)	VOE_s Max (%)
1	L	2.51	3.36±6.57	0.80	58.80
	R	2.84	3.88±8.21	0.84	71.30
2	L	2.90	3.79±5.84	1.11	48.61
	R	3.14	3.99±5.60	1.23	48.17
3	L	2.22	2.42±1.75	0.86	9.89
	R	2.18	2.51±2.42	0.62	19.38
4	L	3.26	3.38±1.77	0.68	9.94
	R	3.54	3.92±1.91	1.75	16.59
5	L	2.41	3.59±8.94	0.74	77.38
	R	2.27	3.12±6.00	0.87	52.47
6	L	2.62	2.76±1.85	1.24	13.41
	R	2.67	2.88±2.18	1.28	16.06
Avg.		2.71±0.44	3.30±0.56	1.00±0.33	36.84±25.17

The precision of our extracted contours is displayed in Table 4.9 with Symmetric Surface Distances (SSD). This table includes the global average ($SSD_{g.avg}$), an average of slice averages ($SSD_{s.avg Avg}$) and a maximal surface distance (SSD_{max}).

4.6 Discussion

Table 4.9 shows that our segmentation strategy returns average results which are quite close to the ground truth. The vast majority of our contours are within 0-1 voxels of the ground truth's contour. This is within the expected variability typical for manual segmentations (1-2 voxels) (Kaus, Pekar, Lorenz, Truyen, Lobregt, & Weese, 2003).

The last column of Table 4.9 highlights a model's least accurate contour. Across each one of our models, these least accurate contours could be counted to be in one of three locations on the femur (a) the crown of the femur head (b)

the slope of the femur neck directly below the head and (c) the narrow concavity between the greater trochanter and the femur neck.

Table 4.9 Symmetric Surface Distances

ID #		$SSD_{g.avg} \pm$ stdev (mm)	$SSD_{s.avg} Avg. \pm$ stdev (mm)	SSD_{max} (mm)
1	L	0.26±0.42	0.29±0.27	4.83
	R	0.29±0.46	0.33±0.44	6.30
2	L	0.28±0.41	0.31±0.25	3.49
	R	0.30±0.44	0.32±0.22	4.32
3	L	0.23±0.38	0.22±0.11	2.76
	R	0.22±0.37	0.22±0.11	2.76
4	L	0.31±0.43	0.30±0.12	3.42
	R	0.37±0.40	0.37±0.09	3.16
5	L	0.26±0.45	0.31±0.51	7.36
	R	0.25±0.42	0.28±0.33	4.93
6	L	0.29±0.38	0.29±0.13	2.16
	R	0.29±0.39	0.29±0.14	2.50
Avg.		0.28±0.04	0.29±0.04	4.00±1.60

Errors in the first two locations can be visualized in the profiles of femur models obtained through our method, seen in Table 4.6 and Table 4.7. We attribute errors in these two regions to blurred and noisy cross-sections of the bone surfaces.

Errors, in the gap between the greater trochanter and the femur neck usually occur when the gap is particularly narrow and deep. This issue is native to most snake segmentation schemes and can be attributed to the external forces acting on the evolving contour cancelling each other out when inside narrow concavities (Xu & Prince, 1998). Possible fixes to this shortcoming include increasing the upscale factor during the pre-processing phase or to alter the external force model to better carry the flow of gradients.

The most recent femur-segmentation strategy against which we can compare our quantitative results is (Schmid, Kim, & Magnenat-Thalman, 2011). Schmid et al. reported and an $SSD_{g.avg}$ of 1.21 ± 0.53 mm, a SSD_{max} of 7.57 ± 2.46 and a VOE_g of 18.02 ± 6.12 % with their high-resolution, low field-of-view, MRI scans of the human femurs. The high-field of view of our scans is a partial contributor to our improvement over these results.

Generally, our segmentation results for the femur body better matched the manual segmentation than the results from the femur head. This is especially true for the top-most slices of the femur head which are the leading location of our SSD_{max} . This is recognized as being due to the tubular regions of the femur having a much thicker shell of compact bone than can be found around the femur head's. This supplies the snake algorithm with strong stopping criteria which halts the curve's evolution.

5 Conclusion

In this paper, we described a method to segment a femur from a CT scan. We initiate the segmentation by having the user subdivide the femur into two regions (the near-spherical head and the vertically tubular body) and used two levels of segmentation (rough and fine). A number of pre-processing steps are employed on each ROI's volume to improve the shape and accuracy of our methodology, such as replacing

sections of the acetabulum with low-gradient fill. Afterwards, we segment the femur-head followed by the femur body using a sequential set of related morphological snake operations.

We find the qualitative results returned by our method to be extremely encouraging given the diverse set of bones on which we experimented. We were able to return results which were largely within the expected range of error for manual segmentations. The user initialization can take under a minute to perform while the segmentation process usually took roughly 5-10 minutes to complete on a modern desktop computer, which is a staggering improvement over the hours it might take a radiologist to manually segment a femur.

The next step in our research is to implement a 3D implementation of morphological snakes to see if we can improve results on error-prone locations on the femur.

6 References

- Alvarez, L. B.-N. (2010). *A Real Time Morphological Snakes Algorithm*. Retrieved 06 09, 2010, from Image Processing On Line : http://www.ipol.im/pub/algo/abmh_real_time_morphological_snakes_algorithm/
- Álvarez, L., Baumela, L., Henríquez, P., & Márquez-Neila, P. (2010). Morphological Snakes. *Computer Vision and Pattern Recognition, IEEE Conference on*. San Francisco.
- Caselles, V., Kimmel, R., & Sapiro, G. (1997). Geodesic Active Contours. *Computer Vision, International Journal of*, 22(1), 61-79.
- Chan, T., & Vese, L. A. (2001). Active Contours Without Edges. *Images Processing, IEEE Transactions on*, 10(2), 266-277.
- Chen, S., & Radke, R. J. (2009). Level Set Segmentation with Both Shape and Intensity Priors. *Computer Vision, IEEE 12th International Conference on*, (pp. 763-770).
- Cohen, L. (1991). On Active Contour Models and balloons. *CVGIP: Image Understanding*, 53(2), 211-218.
- Cohen, L. D., & Cohen, I. (1993). Finite-Element Methods for Active Contour Models and Balloons for 2-D and 3-D Images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 15(11), 1131-1147.
- Cootes, T. F., Edwards, G. J., & Taylor, C. J. (1998). Active Appearance Models. *European Conference on Computer Vision, Proceedings on*, 2, 484-4498.
- Cootes, T. F., Taylor, C. J., Cooper, D. H., & Graham, J. (1995). Active Shape Models - Their Training and Application. *Computer Vision and Image Understanding*, 61(1), 38-59.
- Cremers, D., Osher, S. J., & Soatto, S. (2006). Kernel Density Estimation and Intrinsic Alignment for Shape Priors in Level Set Segmentation. *Computer Vision, International Journal of*, 69(3), 335-351.
- Cremers, D., Schnörr, C., Weickert, J., & Schellewald, C. (2000). Diffusion-Snakes Using Statistical Shape Knowledge. *Notes in Computer Science*, 1888, 164-174.
- Fernandes, L. A., & Oliveira, M. M. (2008, January). Real-time detection through an improved Hough transform voting scheme. *Pattern Recognition*, 41(1), 299-314.
- Field, D., & Hutchinson, J. O. (2008). *Field's Lower Limb Anatomy, Palpation and Surface Markings* (1st ed.). Churchill Livingstone.
- Gilles, B., Moccozet, L., & Magnenat-Thalmann, N. (2006). Anatomical Modelling of the Musculoskeletal Systems. *Medical Image Computing and Computer Assisted Intervention, International Conference on*, (pp. 289-296).
- He, L., Peng, Z., Everding, B., Wang, X., Han, C. Y., Weiss, K. L., et al. (2008). A comparative study of deformable contour methods on medical image segmentation. *Image and Vision Computing*, 141-163.
- Hossain, M., & Andrew, J. G. (2008). Current Management of femoro-acetabular impingement. *Current Orthopaedics*, 22(4), 300-310.
- Kass, K., Witkin, A., & Terzopoulos, D. (1987). Snakes: Active Contour Models. *International Journal of Computer Vision*, 1(4), 321-331.
- Kaus, M. R., Pekar, V., Lorenz, C., Truyen, R., Lobregt, S., & Weese, J. (2003). Automated 3-D PDM Construction From Segmented Images Using Deformable Models. *Medical Imaging, IEEE Transactions on*, 22(8), 1005-1013.
- Magnenat-Thalmann, N., Yahia-Cherif, L., Gilles, B., & Molet, T. (2003). Hip Joint Reconstruction and Motion Visualization Using MRI and Optical Motion Capture. *Biomedizinische Technik/Biomedical Engineering*, 48(s1), 20-23.
- McInerney, T., & Terzopoulos, D. (2000). T-snakes: Topology Adaptive Snakes. *Medical Image Analysis*, 4(2), 73-91.
- Rizzo, D. C. (2001). *Delmar's Fundamentals of Anatomy & Physiology*. Albany, United States of America: Delmar a division of Thompson Learning.
- Schmid, J., Kim, J., & Magnenat-Thalman, N. (2011). Robust statistical shape models for MRI bone segmentation in presence of small field of view. *Medical Image Analysis*, 15, 155-168.
- Sherouse, G. W., Novins, K., & Chaney, E. L. (1990, March). Computation of digitally reconstructed radiographs for use in radiotherapy treatment design. *Radiation Oncology * Biology * Physics, International Journal of*, 18(3), 651-658.
- Song, W. W., Li, G. H., Ou, Z. Y., Han, J., Zhao, D. W., & Wang, W. M. (2007). Model-Based Segmentation of Femoral Head and Acetabulum from CT Images. *Complex Medical Engineering, IEEE/ICME International Conference on*, 586-590.
- Stegmaier, S., Strengert, M., Klein, T., & Ertl, T. (2005). A Simple and Flexible Volume Rendering Framework for Graphics-Hardware-based Raycasting. *Volume Graphics, Proceedings of*, (pp. 187-195). New York, USA.
- Tannast, M., Siebenrock, K. A., & Anderson, S. E. (2007). Femoroacetabular Impingement: Radiographic Diagnosis - What the Radiologist Should Know. *American Journal of Roentgenology*, 188, 1540-1552.
- Vese, L. A., & Chan, T. F. (2002). A Multiphase Level Set Framework for Image Segmentation Using the Mumford and Shah Model. *Computer Vision, International Journal of*, 50(3), 271-293.

- Xu, C., & Prince, J. L. (1998). Snakes, Shapes, and Gradient Vector Flow. *Image Processing, IEEE Transactions on*, 7(3), 359-369.
- Yokota, F., Okada, T., Takao, M., Sugano, N., Tada, Y., & Sato, Y. (2009). Automated Segmentation of the Femur and Pelvis from 3D CT Data of Diseased Hip Using Hierarchical Statistical Shape Model of Joint Structure. *Medical Image Computing and Computer-Assisted Intervention: Lecture Notes from Computer Science*, 5762, 811-818.
- Zoroofi, R. A., Sato, T., Sasama, T., Nishii, T., Nobuhiko, S., Yonenobu, K., et al. (2003). Automated Segmentation of Acetabulum and Femoral Head From 3-D CT Images. *Information Technology in Biomedicine, IEEE Transactions on*, 7(4), 329-343.