Efficient Segmentation of 3D Fluoroscopic Datasets from Mobile C-arm

Martin Styner, Haydar Talib, Digvijay Singh, Lutz-Peter Nolte

M.E. Müller Research Center for Orthopaedic Surgery, Institute for Surgical Technology and Biomechanics, University of Bern, P.O. Box 8354, 3001 Bern

ABSTRACT

The emerging mobile fluoroscopic 3D technology linked with a navigation system combines the advantages of CT-based and C-arm-based navigation. The intra-operative, automatic segmentation of 3D fluoroscopy datasets enables the combined visualization of surgical instruments and anatomical structures for enhanced planning, surgical eye-navigation and landmark digitization. We performed a thorough evaluation of several segmentation algorithms using a large set of data from different anatomical regions and man-made phantom objects. The analyzed segmentation methods include automatic thresholding, morphological operations, an adapted region growing method and an implicit 3D geodesic snake method. In regard to computational efficiency, all methods performed within acceptable limits on a standard Desktop PC (30sec-5min). In general, the best results were obtained with datasets from long bones, followed by extremities. The segmentations of spine, pelvis and shoulder datasets were generally of poorer quality. As expected, the threshold-based methods produced the worst results. The combined thresholding and morphological operations methods were considered appropriate for a smaller set of clean images. The region growing method performed generally much better in regard to computational efficiency and segmentation correctness, especially for datasets of joints, and lumbar and cervical spine regions. The less efficient implicit snake method was able to additionally remove wrongly segmented skin tissue regions. This study presents a step towards efficient intra-operative segmentation of 3D fluoroscopy datasets, but there is room for improvement. Next, we plan to study model-based approaches for datasets from the knee and hip joint region, which would be thenceforth applied to all anatomical regions in our continuing development of an ideal segmentation procedure for 3D fluoroscopic images.

Keywords: Segmentation, Cone beam CT, Fluoroscopy, Region Growing, Geodesic Snake, Orthopaedic surgery, Computer Assisted Surgery

1. INTRODUCTION

Today's routine imaging technologies in the operating room consist mostly of 2D imagery based on ultrasound and fluoroscopy. Intra-operative 3D imaging modalities based on Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) are in development, but are far too expensive for routine use in most hospitals. The emerging mobile fluoroscopic 3D technology (C-arm equipment based) is expected to fill this imaging gap in orthopaedic surgery.[1–3] The availability of 3D imaging in the OR provides full information about the current anatomical setting. Movement of bone structures or repositioning can be imaged as they are and are not influenced by patient positioning or surgical treatment. Due to new developments, which further enhance the imaging quality, mobile C-arm technology can also be used in soft-tissue surgeries.[4]

In Computer Assisted Surgery (CAS) 3D fluoroscopy linked with a navigation system combines the advantages of CT-based and C-arm-based navigation. The use of an additional navigation system creates an inherent registration with the 3D image, such that surgical instruments can be visualized along with the current patient anatomy.[5, 6] The CAS systems linking navigation and 3D fluoroscopy visualize, thus far, only orthogonal and oblique slices through the 3D image. This visualization is appropriate for most surgeries of smaller joints such as the knee, elbow, hand and foot. In spine, shoulder, hip and pelvic surgeries, the surgical hand-eye coordination and anatomical orientation are often impaired due to the degraded image quality and the small field-of-view of

Email: martin_styner@ieee.org, Haydar.Talib@MEMcenter.unibe.ch, Digvijay.Singh@MEMcenter.unibe.ch, Lutz.Nolte@MEMcenter.unibe.ch. WWW: www.memcenter.unibe.ch
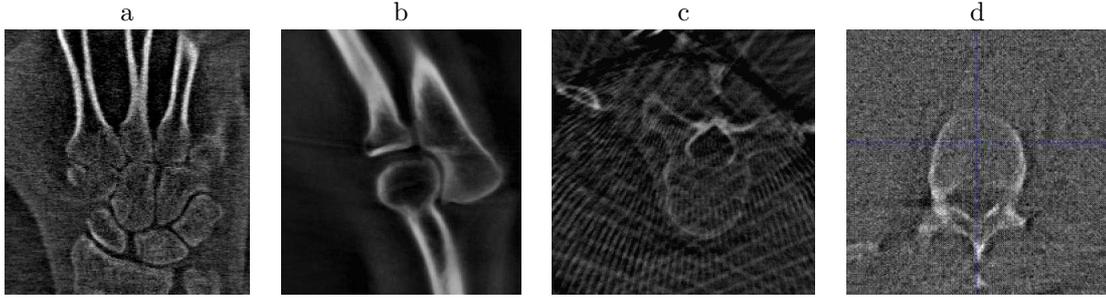
Figure 1. Six Exemplary images from mobile iso-centric C-arm equipment. Relatively good quality in images of the extremities such as the hand(a) and the elbow (b). The effects of low contrast, artifact formation and high noise level on the image quality are likely to be observed in images of the spine (c, d)

12.8cm^3. The intra-operative, automatic segmentation of 3D fluoroscopy datasets would enable the combined 3D-rendered visualization of surgical instruments and anatomical structures. This visualization is likely to enhance the intra-operative planning step, surgical eye-navigation and the digitization of anatomical landmarks.

Segmentation methods that were developed for CT or MR datasets cannot be directly applied to 3D fluoroscopy datasets due to specific artifacts; generated mainly by the low number of projections used for the reconstruction as described in section 2.1. Currently semi-automatic thresholding followed by connectivity-based component selection is mainly used for segmenting 3D fluoroscopy datasets. This method is unfit for intra-operative use since it demands a high degree of user interaction and the resulting segmentation is usually of rather bad quality. Since little is known about the segmentation of 3D fluoroscopy datasets, we performed a thorough evaluation of several adapted standard algorithms using a large set of cadaveric and patient data. Considering that 3D fluoroscopy is employed in trauma surgeries of fractured bone structures, all of the studied methods are intensity-based and do not employ a statistical model-based technique. For specific applications, the use of model-based methods would be advantageous and will be part of our future research.

Sections 2.1 and 2.2 describe the 3D fluoroscopic-image properties and present the image database used in all tests, respectively. The remainder of section 2 describes the methods that were applied to the image database. Section 3 presents the results of the performance of the segmentation methods and compares them with one another in regard to anatomical correctness, computational efficiency and usability for intra-operative visualization.

2. METHODS

Due to the limited knowledge in published literature about the segmentation of 3D fluoroscopy datasets, we performed a thorough evaluation of several adapted standard algorithms using a large set of cadaveric and patient data. One of the intra-operative requirements for the use of segmentation algorithms is the minimal amount of user interaction. Thus, we first studied automatic thresholding algorithms based on Otsu's method and on histogram model fitting. Following the threshold computation, morphological operators along with connectivity operators were investigated, as well as an adapted region growing method, described in sections 2.3 and 2.5. The region growing method was also used as initialization for a geodesic implicit deformable surface segmentation.

2.1. 3D Volumetric Images from Mobile Iso-centric C-arm Fluoroscopy Equipment

The images studied in this paper were acquired with a standard Siemens Siremobil IsoC^{3D} scanner, which combines both iso-centric C-arm 2D projective fluoroscopy and 3D cone-beam reconstruction image acquisition. In 3D fluoroscopy a set of 2D X-ray beam projections is captured on a circular arc by an image intensifier tube. The 3D volumetric images have a small field-of-view of 12.8cm^3 and a high resolution with a voxel size of 0.5mm^3. The reconstructed images often suffer from an abundance of noise and artifacts in comparison with standard CT imagery. There are several reasons for this finding: 1) The circular arc covers at most 190 degrees
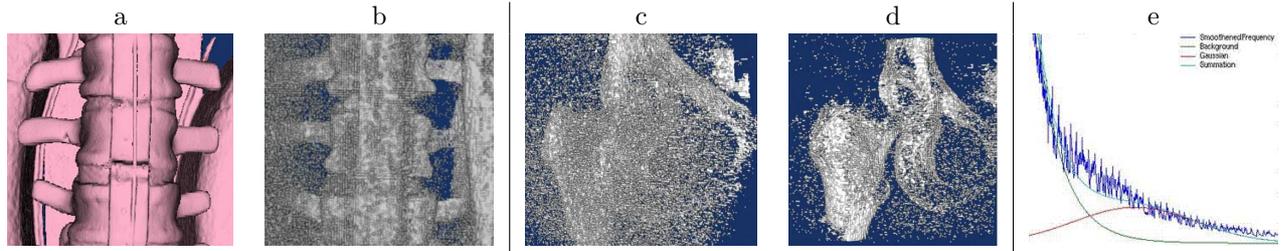
Figure 2. Thresholding: (a+b) Rendering of optimal thresholding applied to a CT (a) and 3D fluoroscopy (b) dataset of a plastic spine. (c+d) Rendering of thresholding result of a hip dataset with the optimal threshold (c) and a slightly higher threshold (d). Neither case results in a satisfying segmentation. (e) Bone section of the histogram from a 3D fluoroscopic image of good quality. It is hard to clearly detect the bone class from the histogram

instead of a full 360 degrees resulting in a set of artifacts. 2) The X-ray beam is detected by an image intensifier tube, which gives rise to a higher level of noise and introduces geometric distortions that need to be corrected. 3) The projections are acquired less densely-spaced, causing line-shaped artifacts (see Fig. 1c). Nevertheless human experts can efficiently interpret 3D fluoroscopic images. The use of this technology showcases advancement for routine intra-operative imaging; especially for surgeries in extremities. Due to the degraded image quality and the limited field-of-view in spine and pelvic surgeries, however, a segmentation algorithm, even if only partially successful, would result in an enhanced understanding of the anatomy in a given image.

We investigated the use of noise reduction to improve image quality. Since 3D fluoroscopy is used intra-operatively, most sophisticated methods such as shape-based, or diffusion methods are not computationally efficient enough. Simple methods like Gaussian, mean, median or morphological smoothing remove many edges that are weak due to the low contrast. We observed that noise reduction is only an option in images with good contrast that are not in need of noise reduction. Thus, we did not employ noise reduction in our studies.

2.2. Database of 3D Fluoroscopy Images

We collected a database of 3D fluoroscopic images in order to evaluate the performance and the parameter-settings of the studied segmentation methods. We acquired images from man-made objects of fully known geometry, as well as testing data from plastic and cadaveric bones, and also of clinical patient data. The patient and cadaver data was acquired from different anatomical regions (hand, elbow, shoulder, head, spine, femur, knee, tibia, foot) focusing on bone structures, totaling over 80 datasets. The man-made aluminum objects datasets consisted of full and hollow cylinders, spheres of different size, as well as two thin plates. Additionally CT and 3D fluoroscopic images were acquired from plastic bones using the CT as ground truth. Since there are no motion artifacts, few line-shaped artifacts and a very high contrast between air and the plastic bone, the image quality of these datasets can be regarded as the best possible. Even in this case the 3D fluoroscopic images have considerable noise and degraded edges as can be seen in Fig. 2b, which shows a rendering of the result from a manually optimized thresholding.

2.3. Automatic Thresholding: Otsu

The challenge of segmenting medical images has often been approached with the use of thresholding, particularly of Otsu's optimal threshold selection method.[7] The method performs very well when trying to separate one foreground object from the background. More specifically, Otsu's method intends to separate two distinct histogram classes in a given image by maximizing the between-class variance over the entire histogram.

As an initial step Otsu's algorithm appeared to be a sound choice in attempting to segment 3D fluoroscopic medical images of bone regions, as the images are expected to contain two histogram classes. In analyzing some preliminary next-generation "Flat Panel"[4] medical images, we encountered a problem typically found in CT images as well, where Otsu's algorithm for the two-class case was no longer valid, as three distinct classes were detectable in the histogram. We then expanded Otsu's original method to the three-class case, which should yield satisfactory results though he warns "that the selected thresholds generally become less credible as the

number of classes to be separated increases".[7] Detailed information regarding Otsu's method can be found in his paper,[7] but we reprint the mathematical expressions to clarify the derivation of the three-class threshold selection. Consider an image with grey levels represented by $[1, 2, \ldots, L]$, with the histogram normalized to give a probability distribution over the image. We then consider separating the histogram into two classes, C_0 and C_1, by a threshold at pixel intensity-level k. Now the image can be represented by two sets; $[1, \ldots, k]$ and $[k+1, \ldots, L]$. Otsu defines the probabilities of class occurrences and the class mean levels, which are respectively given by:

$$\omega_0 = \sum_{i=1}^{k} p_i, \quad \omega_1 = \sum_{i=k+1}^{L} p_i = 1 - \omega_0 \tag{1}$$

$$\mu_0 = \sum_{i=1}^{k} ip_i/\omega_0, \quad \mu_1 = \sum_{i=k+1}^{L} ip_i/\omega_1 = \frac{\mu_T - \omega_0\mu_0}{1 - \omega_0} \tag{2}$$

where μ_T is the total mean level of the original image, and p_i is the probability of occurrence of intensity-level i. Otsu then shows that in order to determine the optimal threshold k to separate the two classes, it is sufficient to maximize the between-class variance, σ_B^2, which is given by:

$$\sigma_B^2 = \omega_0(\mu_0 - \mu_T)^2 + \omega_1(\mu_1 - \mu_T)^2 \tag{3}$$

Now we consider separating the image histogram into three classes: C_0 for $[1, .., k_0]$, C_1 for $[k_0 + 1, .., k_1]$ and C_2 for $[k_1 + 1, .., L]$ such that the two thresholds k_0 and k_1 satisfy $1 \leq k_0 < k_1 < L$. We then define the probabilities of class occurrences and the class mean levels for the three classes:

$$\omega_0 = \sum_{i=1}^{k_0} p_i, \quad \omega_1 = \sum_{i=k_0+1}^{k_1} p_i, \quad \omega_2 = \sum_{i=k_1+1}^{L} p_i = 1 - \omega_1 - \omega_0 \tag{4}$$

$$\mu_0 = \sum_{i=1}^{k_0} ip_i/\omega_0, \quad \mu_1 = \sum_{i=k_0+1}^{k_1} ip_i/\omega_1, \quad \mu_2 = \sum_{i=k_1+1}^{L} ip_i/\omega_2 = \frac{\mu_T - \omega_1\mu_1 - \omega_0\mu_0}{1 - \omega_1 - \omega_0} \tag{5}$$

In accordance with Otsu's derivation, the between-class variance σ_B^2 becomes:

$$\sigma_B^2 = \omega_0(\mu_0 - \mu_T)^2 + \omega_1(\mu_1 - \mu_T)^2 + \omega_2(\mu_2 - \mu_T)^2 \tag{6}$$

2.4. Automatic Thresholding: Histogram Model Fit

Parallel to the Otsu thresholding, we developed a second automatic thresholding method based on the histogram-fitting approach employed for MR images. We first constructed a mathematical class model for the histogram, which is then fitted to the observed-image histogram by minimizing the difference of the two(see Fig. 2e). We studied several models assigning exponential and Gaussian-distributed classes to the background, soft tissue and bone. The background is best modeled as an exponential class and bone as a single Gaussian. Since soft tissue is barely discernable from background in most images, the most stable class model was achieved with a joint exponential background/soft tissue class. Median smoothing is applied to the observed histogram due to the presence of noise and high amplitude spikes.

2.5. Morphological Operations

As we will later discuss in more detail, using Otsu's method for optimal threshold selection for 3D fluoroscopic images seldom yields satisfactory segmentation of bone regions, due to the low-contrast nature of the images. Additionally, we implemented several morphological operations[8, 9]: opening, closing and connected-component

labeling(CCL), the latter to identify disconnected non-bone-regions for removal from the segmentation. Opening uses the morphological operator erosion, followed by a dilation. Closing consists of a dilation followed by erosion.

Opening would be an adequate operation when an image is under-thresholded; that is, when parts of the background are also segmented with the foreground object. Closing would be more appropriate for cases where there is over-thresholding, wherein parts of the foreground object are not included in the segmentation. We used 6-connected CCL to label regions at the final stage of segmentation, to finally eliminate regions that represent less than 1% of the segmented image volume. Our rationale was that low-volume regions were due to remnants of noise or artifacts that were wrongfully segmented.

Combining the morphological operations with the Otsu optimal threshold selection, we implemented three approaches to achieve the intensity-based segmentation of IsoC^{3D} images. The first approach was to threshold using the two-class Otsu, followed by a closing operation. The second approach was to threshold using Otsu, followed by an opening, and finally by a CCL with 1%-removal. The third approach was to threshold using Otsu, followed by an opening, then a closing, and finally by a CCL with 1%-removal.

2.6. Leak-Constrained Region Growing

Other approaches taken to segment medical images involve the use of region-based algorithms, which comprise a method to grow from seed points located in areas relevant for segmentation. A good review is presented in Zucker's survey on region growing.[10]

We developed a region-based segmentation algorithm using connected-threshold region growing at the core. The connected-threshold method grows from seed points by checking neighboring voxels to see if any meet the intensity threshold requirement, which we will refer to as the growing threshold. This method used 6-connectedness neighborhoods. Voxels that have higher intensity than the growing threshold are segmented, and the method continues to grow from them as well, until none of the remaining neighbors meet the growing criterion. It is clear from this description that the method requires two inputs: the threshold value that controls growing, and a list of seed points from which to grow. We attempt to automate these inputs to minimize user-interactivity.

The list of seed points is obtained by over-thresholding the input image, followed by an erosion step, so that we are only left with small areas indicating parts of the bone regions having the greatest intensity. How to determine the threshold value used to obtain the seed list will be discussed shortly.

Since the 3D fluoroscopic images have low bone-to-background contrast, we expanded the given connected-threshold method to allow more freedom when growing in bone structures, while minimizing growth into background areas. A low-contrast image means that the bone regions we are interested in segmenting will not have widely differing intensity values than the background area. As such, a given growing threshold, if set too high, would not segment relevant bone structures of low intensity. Conversely, a low growing threshold would quickly cause the region growing algorithm to flood out into the background. Since we want to visualize as much of the bone region as possible, we needed the freedom to lower the growing threshold such that darker areas of the bone would be included in the segmentation. In doing so, we added two constraints in the connected-threshold method.

The first constraint takes into account edge information in the image. The gradient magnitude is taken over the entire image. As the algorithm performs the neighbor checks, it considers the gradient magnitude in addition to the intensity. If the gradient magnitude of a given neighboring voxel is above a threshold value, indicating an edge, then that neighbor would be included in the segmentation though the algorithm does not grow from that point. Now, with this added gradient constraint, we have a third parameter needed as an input: the gradient magnitude threshold. The automation in selecting the latter will discussed shortly.

The second constraint corrects for leaking into the background. Again, due to the nature of the images, we often encounter the case of not only having low-contrast, but also of weak edges at bone boundaries, wherein the gradient map shows broken contours at the edges. Consequently, with no gradient-stoppage at weak edges, the algorithm floods out into the background. The leak constraint involves eliminating leaks of a specified size, in voxels, by making use of a modified opening operation with an embedded CCL step. As the first step of the opening, the erosion has a diameter equal to the width of leaks we wish to eliminate. Once the leaks themselves are eroded from the segmented image, we would then still be left with areas that were falsely segmented. A
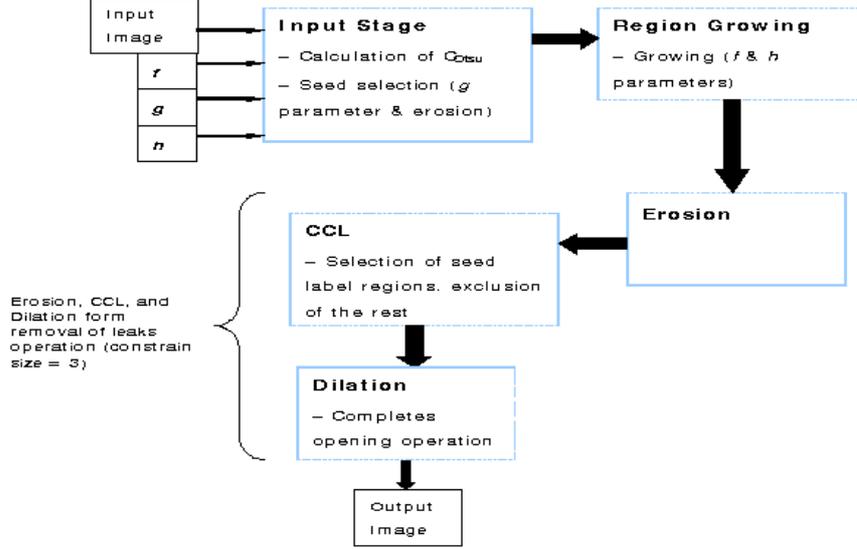
Figure 3. Flow diagram for the region growing system

CCL is then performed, and all regions that do not contain one of the original seed points are removed from the segmentation. A dilation is finally applied to return the segmented image to its original size, completing the opening operation. Now there is a fourth input parameter needed by our region-growing system: the leak-constraining size. The region growing system is shown in Fig. 3 in a flow diagram.

Our region-growing system now needs four input parameters, and we do not want to make them user-defined, as this would prove highly inefficient in computer-assisted surgery. To automate the setting of parameters we base them all, except for the leak-constraining size, on the threshold value computed by Otsu's algorithm, I_{Otsu}. We then define $C_{Otsu} = I_{Otsu} - I_{1\%}$, where $I_{1\%}$ is the intensity value below which 1% of voxels lie, in the image. $I_{1\%}$ is obtained from the cumulative histogram of the input image, and is chosen as such to reduce the influence of intensity outliers. Using C_{Otsu} we define the following threshold criteria:

$$
\begin{aligned}
T_{grow} &= I_{Otsu} - f \cdot C_{Otsu} & (7) \\
T_{seed} &= I_{Otsu} + g \cdot C_{Otsu} & (8) \\
T_{gradient} &= h \cdot C_{Otsu} & (9)
\end{aligned}
$$

T_{grow}, T_{seed} and $T_{gradient}$ are the threshold values used for growing, seed selection and gradient check respectively. f, g and h are now the only variable parameters. The remaining challenge in automating the system at this point is to determine which values of f, g and h are most appropriate for segmenting IsoC^{3D} images. For our experiments and analysis, we tested several values: $f = (0.05, 0.1, 0.2)$, $g = (1.2)$ and $h = (0.6, 0.8)$.

The final parameter to automate was the value of leak-constraining size that would be most ideal for our experiments. The size of leaks to be prevented was set *a priori* to a value of 3, meaning that any part of the segmentation having a width of three voxels or less was going to be removed. We studied higher values for this constraint, and aside from lower computational efficiency, they did not produce good results except for a few cases. Choosing a higher constraining parameter may produce a cleaner image, but ultimately there would be loss of significant structures.

2.7. Implicit Geodesic Deformable Surface Segmentation

Deformable surface segmentation approaches are based either on an explicit (discrete mesh or parametric) surface or an implicit surface. Explicit surfaces have the advantage of tightly controlling the topology, whereas implicit

surfaces offer the possibility for topological changes. In our research we first focused on implicit methods, as 3D fluoroscopy is also used in trauma surgeries with bone fractures. Topological change is necessary in these cases for an appropriate segmentation.

Implicit deformable surface methods are commonly based on level-set evolution techniques that embed the surface in a higher dimension and define the implicit surface as the zero level set. Geodesic Deformable Surfaces are appealing especially for volumetric data processing due to their elegant formulation as nonlinear PDE's. The formalism can be naturally extended from 2-D to higher dimensions. Our studies are based on the SNAP tool,[11] which was originally developed at the University of North Carolina. The geodesic deformable surfaces in SNAP implement several varieties presented in the literature, both boundary-driven and region-based methods. We selected as the most appropriate method for 3D fluoroscopic images the region-based method due to high levels of noisy and artifact-bound edges. The fine parameters were tuned to the individual anatomical regions and the segmentation was initialized using the result of the leak-constrained region growing method described in section 2.6.

3. RESULTS

All the segmentation methods performed within acceptable time limits on a standard Desktop PC. The implicit snake method (5-8 min) was less computationally efficient than the other, less sophisticated methods (1 - 2 min). The best segmentation results were obtained with datasets from long bones (tibia, femur), followed by datasets from extremities (hand, head, ankle, knee, wrist, elbow). The segmentations of spine (cervical, thoracic, lumbar), pelvis and shoulder datasets were generally of poorer quality using the studied methods. As expected, the purely threshold-based methods produced the worst results. Several of the studied combinations of morphological operations and automatic thresholding methods significantly improved those results. The best results were obtained using the adapted region growing method described in section 2.6. This system efficiently removed many wrongly segmented regions that were due to noise and artifacts. These segmentations were considered appropriate for a smaller set of clean images from different anatomical regions, and provided the best results for datasets of joints and in the lumbar and cervical spine areas. The implicit snake method produced more accurate and smoother segmentations and was able to remove wrongly segmented regions originating from skin tissue, which was predominantly present in images of extremities. In regard to the human-made objects, the implicit snake method performed considerably better than the other methods.

3.1. Automatic Thresholding, Morphology and Connectivity Operators

Otsu's method for optimal threshold selection produced adequate values for the cleaner, higher contrast images. The more corrupted images, where a separate foreground class was hard to detect in the histogram, posed a problem for the threshold calculator. The system, however, did not depend solely on the value for the Otsu threshold, and despite a badly-thresholded initial image, the rest of the system proceeded to improve the quality of the segmentation. The histogram matching only worked for a few cases among the cleanest images. Individual datasets' histograms did not contain any visible foreground class, causing this method to fail. Thereby for the remainder of the evaluation we will only discuss the results of the Otsu-based segmentation.

Fig. 4 demonstrates some results of interesting cases for discussion. The elbow scan has a high contrast between the bone region and the background, making segmentation of the joint easier. The picture in Fig. 4b demonstrates a satisfactory result. In this particular case the morphological operations had some very slight, yet observable negative effects upon the Otsu-segmented image. The closing operation(Fig. 4c) seems to have solidified bone boundaries, yet falsely joined the humerus to the ulna. Applying an opening operation(Fig. 4d) thinned the bone boundaries, creating a gap at the edge of the ulna. We can already observe, then, the reasoning behind choosing the sequence of operations as described in 2.3 and 2.5. From this case we can surmise that a small loss from opening following an Otsu-segmentation could be recovered by a closing operation. The case of the lumbar spine demonstrates how the operations are affected by noise at different stages. Though the bone region is still visible to the human eye after an Otsu-segmentation(Fig. 4f), the resulting image is nonetheless highly corrupted by noise. A closing step(Fig. 4g) worsens the segmented image, as the speckles of noise become large blobs that fill most of the image space, making it impossible to see the bone region. An opening(Fig. 4h) operation, however, produces a radical improvement in cleaning up the noise, with some loss to bone structures.
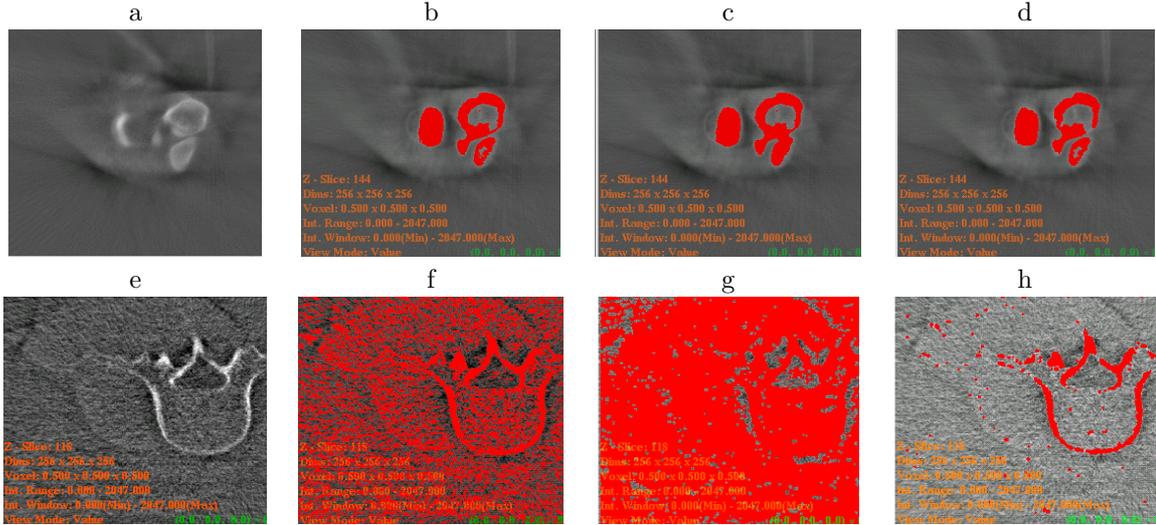
Figure 4. (a) Cadaver elbow input image, (b) Otsu threshold applied to *(a)*, (c) Closing applied to *(b)*, (d) Opening applied to *(b)*, (e) Patient lumbar input input image, (f) Otsu threshold applied to *(e)*, (g) Closing applied to *(f)*, (d) Opening applied to *(f)*

In many cases, after opening and closing operations, applied in sequence to Otsu-segmentations, there still remained non-bone traces in the output image. To finally eliminate these erroneous areas, the CCL was applied as a final stage. Though the CCL identified noise and small-scale artifacts, it sometimes removed bone structures as well.

At the end of the analysis, we were able to make the following observations:

Otsu Thresholding: This produced a satisfactory result when the input image was fairly clean, and had high-contrast to begin with. Often, though, more post-processing was required to achieve better segmentation.

Opening: Of the morphological operations this is the best overall choice as a second step in the segmentation.

Closing: This proved to be a poor second step operation, as it amplified the effects of artifacts and noise in the image, and usually degrades the output quality of the Otsu-segmentation stage.

Otsu+Opening+Closing: There are some cases when the closing step degraded quality after the opening stage, but this occurred when the image was very poor to begin with. In most cases, the closing step improved image quality. Closing was a well-chosen addition, since it did little in terms of amplifying the effects of noise once the opening had been applied. The closing also served to join relevant fragmented boundaries or structures.

Otsu+Opening+CCL: This sequence cleaned out noise in many cases, but removed relevant structures.

Otsu+Opening+Closing+CCL: Although this method was slightly less successful in removing noise than the previous case, it retained important structures more frequently.

As a final note concerning this approach to segmenting 3D fluoroscopic image data, we will briefly discuss the influence of the different factors to the input image quality. When noise was the main contributing factor to poor images, a dramatic increase in segmentation quality was observed at the opening stage. This suggests that noise, to a certain degree, is the more desirable corrupting factor to have in the image scans. The presence of artifacts in a given scan often cause poor segmentations.

The conclusion of this approach is that intensity-based methods will often fail to produce adequate segmentations. In our analysis we observed a tradeoff, in that to achieve cleaner segmentations, relevant structures were often removed. This tradeoff was observed at the closing stage, when used as a third step in segmentation. Closing prevented the removal of some trace noise, but at the boon of retaining pertinent bone structures in the segmented image. We judged that with medical images it is preferable to have less loss of bone structure rather
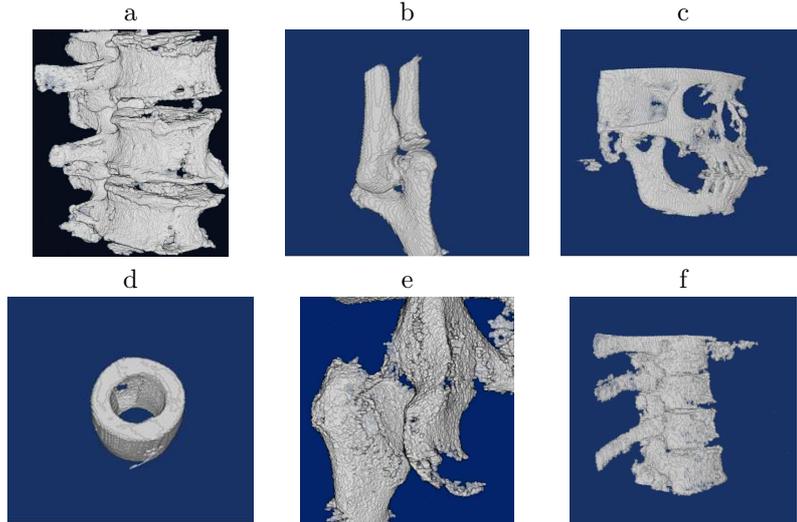
Figure 5. Region growing segmentation of (a) cadaveric thoracic spine, (b) cadaveric elbow, (c) cadaveric cranium (d) hollow man-made cylinder (e) cadaveric pelvis (f) patient thoracic spine

than a cleaner image with loss, since what is lost is not easily recoverable. Since the intensity-based techniques remove structures in some of the cleaner segmentations, their intra-operative use would be limited.

The three-class Otsu threshold selection was studied in some test cases of CT and "Flat Panel" images, as well as some phantom tests. For the phantom tests, the method correctly separated two foreground objects of different intensities from the background as well as from one another. When objects begin to have similar intensity in a given image then the method fails to separate them. An alternative approach was taken to separate three classes, to see if the algorithm performed normally. A normal, two-class Otsu threshold selection was used twice in sequence on an image histogram. This approach behaved slightly better in the observed cases, as the second threshold selection was often more appropriate. The correctness of the obtained Otsu thresholds was qualitatively judged from the histogram.

3.2. Region Growing

Our region growing method was applied and analyzed for the same datasets as for our evaluation in section 3.1. We also propose a similar, albeit more detailed qualitative evaluation. Different values were tested for the f, g and h parameters, as discussed in section 2.6. The single value for the g was deemed suitable for seed point selection in most cases. The parameters that had strongest influence on the segmentation process were the f and h parameters, which control the flooding and edge-stopping respectively. The higher the f parameter is set, the more growing is allowed, and similarly for the h parameter. For this experiment, the ideal case of segmentation is when the f parameter is set high enough to allow for free growing within seed areas, and that the h parameter is low enough to restrict growth past, yet not within, bone regions.

Fig.5 shows some of the better results of segmentation using this method. Particularly interesting results are those of the thoracic spine, although not ideal segmentations, show that our region growing method is already a better choice than the intensity-based method we examined. Table 1 shows the qualitative evaluation of this algorithm's performance. Though it may seem from the table that much of a particular structure has been included in a segmentation, a low *Noise, Artifacts and Tissues* rating would imply that those regions were falsely included as well. In some cases human tissue was segmented in addition to bone, despite the absence of noise or artifacts, as seen in Fig. 6b.

Our observations from the qualitative analysis, regarding the f and h parameters were:

Gradient, h parameter: In several cases, due to setting the h parameter to either 0.6 or 0.8, restrictions were imposed on growing within structures, worsening the overall segmentation. That is, the gradient threshold

| Input Image Name | Input Image Quality | Ratings | | Parameter Settings |
| | Noise / Artifacts | S (1-5) / NAT (1-10) | | f, g, and h |
		S	NAT	
Man-Made Object	L/M	4.7	8.1	f=0.2, g=1.2, h=0.8
Head	L/M	4.4	6.3	f=0.05, g=1.2, h=0.6 or 0.8
Neck	H/L	5	3.25	f=0.2, g=1.2, h=0.6 or 0.8
Cervical Spine	L/M-H	4.1	2.2	f=0.05, g=1.2, h=0.8
Thoracic Spine	M-H/ M-H	4.2	3.1	f=0.05, g=1.2, h=0.8
Lumbar Spine	M/M-H	3.4	6.4	f=0.05, g=1.2, h=0.8
Shoulder	L/H	4	2.5	f=0.05, g=1.2, h=0.8
Elbow	L/M-H	4.9	7	f=0.2, g=1.2, h=0.8
Wrist	L/H	4.75	3.5	f=0.05, g=1.2, h=0.6 or 0.8
Hand	L-M/M-H	4.7	3	f=0.2, g=1.2, h=0.8
Pelvis	H/M-H	4.6	2.4	f=0.2, g=1.2, h=0.8
Femur	L-M/L	4.1	8.1	f=0.05, g=1.2, h=0.8
Knee	L-M/M	3.8	4	f=0.05, g=1.2, h=0.8
Tibia	L/M	4.6	8.2	f=0.05, g=1.2, h=0.6 or 0.8
Foot	L/M-H	4.6	3.1	f=0.05, g=1.2, h=0.8

Legend
Input Image Quality: Two parameters qualitatively measure presence of noise and artifacts, evaluated as *Low, Medium and High (L, M, H)*.
Structures (1 to 5, in S column): Quantifies presence of relevant bone structure in the segmentation.

- *1*: Heavy loss of structure, with little to no relevant bone structures retained.

- *5*: The ideal state, where bone has been fully segmented with no elements missing.

Noise, Artifacts and Tissue (1 to 10, in NAT column): Quantifies noise, artifacts and/or tissue present in the segmentation.

- *1-3*: Poor quality image. At *3*, artifacts and noise severely obscure the bone image, although some forms can still be recognized. From *1-2*, the image is mostly corrupted by noise and artifacts, and no pertinent forms can be recognized.

- *10*: The ideal state, without visible noise, artifacts or tissue.

Table 1. Qualitative results of the adapted region growing algorithm

parameter was too restrictive on growth within bone regions. In other cases the value of the h parameter appeared to have caused little or no effect on the growing or the leak prevention. We propose the following explanation: high or low intensity noise and artifacts will produce high gradient values and if this occurs on the bone structure than these will hinder growing within that structure. And once the h parameter is set slightly higher to prevent the aforementioned restriction, then there would be more flooding out into the background region, since there is more space for growing in the low-intensity region. The gradient parameter, ultimately, is more restrictive in bone regions than in the background.

Region Growing, f parameter: The changing f parameter creates a tradeoff in most cases of the 3D fluoroscopic images: the more freedom is given to the growing, the more noise is segmented; conversely the more restriction is enforced, the less pertinent structures are retained even though the effects of noise would be reduced. When the image is clear to begin with, then keeping the f parameter high allows for an ideal segmentation.

3.3. Implicit Geodesic Deformable Surface Segmentation

The result of the deformable surface segmentation is strongly influenced by its initialization. When initialized with a set of seed points similar to the region growing method, the resulting segmentation was highly deficient in comparison. When initialized with a rather detailed and good estimate of the surface the method performed well. Its use is therefore constrained to being an improvement upon a prior segmentation step. In our tests, we initialized the implicit surface with the outcome of the our adapted region growing segmentation.

In cases where the region growing segmentation failed (e.g. shoulder), we did not observe any improvement, as expected. For several fluoroscopic images of the extremities, the region growing segmentation included skin tissue in addition to the correct bone segmentation. In those cases, the smoothness/curvature term of the deformable surface segmentation resulted in the partial removal of thin skin layer parts (see also Fig. 6 for the rarely-encountered optimal case of full removal).

For large datasets the computational efficiency of the method is unfit for an intra-operative setting (10-15 minutes on a standard desktop PC). In order to enhance the computational efficiency the method was only applied to the slightly enlarged bounding box of the initialization. The computational efficiency improved to 5-8 minutes. But this is still too high, especially when considering that this method is used for touching up the prior region growing segmentation.
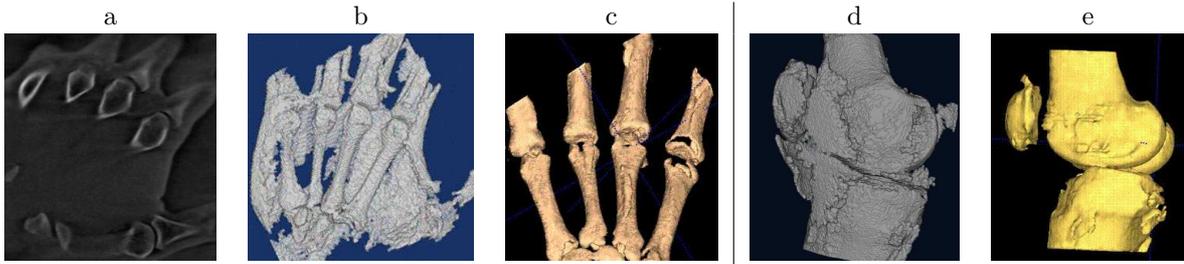
Figure 6. Successful segmentation of a hand (a-c) and knee (d,e) datasets: (a) A rather high contrast and intensity of the skin is present in the 3D fluoroscopic image. (b,d): Region growing segmentation including adjacent skin regions. (c,e): Implicit deformable surface segmentation with full removal of skin regions

4. DISCUSSION

From our intensity-based approach, we observed the impact on segmentation due to the low-contrast nature of the 3D fluoroscopic images. The region growing approach demonstrated that the low contrast also made it difficult to combine leak reduction using gradient edges in hope of liberal growing within bone regions. We did conclude, however, that the region growing method is superior to the intensity-based approach method we used, even though the it is not yet suitable for segmenting more corrupt images, as typically found in the torso or spine areas. It has proved, however, to be computationally efficient and ready to use for anatomical extremities.

The implicit deformable model segmentation is, due to its low computational efficiency and limited scope of application, unfit for a generic intra-operative application. An explicit statistical model segmentation based on models specifically designed for the knee, hip and elbow region has a greater potential for clinical applications.

The results of our experiments fit well with the studies of Rock et al,[12] which show adequate image quality for smaller joints and highly inferior quality for other anatomical regions. For the cranium, hip, lumbar and thoracic spine, the image quality can become too low for adequate intra-operative guidance solely based on the multi-planar slice-view. In these cases, our segmentations visualized in a fully interactive 3D rendering improves the surgical navigation. We expect that statistical model-based segmentation methods, which are currently under development at our institute, will further enhance the segmentation quality. This will lead to additional clinical applications that can benefit from intra-operative 3D fluoroscopy.

5. CONCLUSIONS

We presented our analysis of different methods for bone segmentations from 3D fluoroscopy datasets acquired from a variety of anatomical regions. Our adapted region growing method showed good computational efficiency and segmentation correctness for a large set of images. The implicit snake method provided the best results in regard to segmentation correctness, but was less computationally efficient.

The region growing method performed sufficiently well for extremities such that we plan its use in specific clinical applications for enhancement of anatomical orientation for the surgeon.

This study presents a step forward towards efficient and correct intra-operative segmentation of 3D fluoroscopy datasets, but there is room for improvement. As a next step we plan to study a model-based approach for datasets from selected anatomical regions (spine and hip). The selected approach is known to perform well even in the presence of high noise and artifacts,[13] which we would then apply to all anatomical regions in our continuing development of an ideal segmentation procedure for 3D fluoroscopic images.

ACKNOWLEDGMENTS

REFERENCES

1. M. Kfuri, D. Kendoff, T. Gösling, B. Thumes, T. Hüfner, and C. Krettek, "Iso-c3d control by calcaneus osteosynthesis," in *Computer Assisted Orthopaedic Surgery*, pp. 180–181, Steinkopf, June 2003.

2. E. Euler, S. Wirth, U. Linsenmaier, W. Mutschler, K. Pfeifer, and A. Hebecker, "Comparative study of the quality of c-arm based 3d imaging of the talus," *Unfallchirurg* **104**, pp. 839–46, Sep 2001.

3. M. Richter, D. Kendoff, T. Gösling, J. Geerling, T. Hüfner, and C. Krettek, "Intraoperative 3-d imaging with a mobile image amplifier (iso-c3d) in foot and ankle trauma care," in *Computer Assisted Orthopaedic Surgery*, pp. 302–303, Steinkopf, June 2003.

4. D. Ritter, M. Mitschke, and R. Graumann, "Intraoperative soft tissue 3d reconstruction with a mobile c-arm," in *CARS*, *International Congress Series*, pp. 200–206, Elsevier, March 2003.

5. A.Schmidt, P. Gruetzner, R. Simon, and A. Wentzensen, "Intraoperative 3d imaging in displayed intra-articular calcaneal fractures," in *Computer Assisted Orthopaedic Surgery*, pp. 326–327, Steinkopf, June 2003.

6. P. Grützner, H. Wälti, B. Vock, A. Wentzensen, and L. Nolte, "Inherent spinal navigation using fluoro-ct technology," in *Computer Assisted Orthopaedic Surgery*, pp. 128–129, Steinkopf, June 2003.

7. N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-9**, pp. 62–6, Jan 1979.

8. J. C. Russ, *The Image Processing Handbook*, CRC Press, 4th ed., 2002.

9. M. Sonka, V. Hlavac, and R. Boyle, *Image Processing Analysis and Machine Vision*, PWS Publishing, 2nd ed., 1999.

10. S. W. Zucker, "Region growing: Childhood and adolescence," *Computer Graphics and Image Processing* **5**, pp. 382–99, 1976.

11. S. Ho, H. Cody, and G. Gerig, "Snap: A software package for user-guided geodesic snake segmentation," tech. rep., Dept. Computer Science, University of North Carolina, 2003.

12. C. Rock, D. Kotsianos, U. Linsenmaier, T. Fischer, R. Brandl, F. Vill, S. Wirth, R. Kaltschmidt, E. Euler, K. Pfeifer, and M. Reiser, "Studies on image quality, high contrast resolution and dose for the axial skeleton and limbs with a new, dedicated ct system (iso-c-3 d)," *Rofo Fortschr Geb Rontgenstr Neuen Bildgeb Verfahr.* **2**, pp. 170–6, Feb 2002.

13. D. Shen, E. Herskovits, and C. Davatzikos, "An adaptive focus statistical shape model for segmentation and shape modeling of 3-d brain structures," *IEEE Transactions on Medical Imaging* **20**, pp. 257–270, April 2001.