

# Single click volumetric segmentation of abdominal organs in Computed Tomography images

Brian W. Whitney<sup>a</sup>, Nathan J. Backman<sup>b</sup>, Jacob D. Furst<sup>c</sup>, Daniela S. Raicu<sup>c</sup>

<sup>a</sup>Northern Kentucky University, Nunn Drive, Highland Heights, KY, 41099, USA

<sup>b</sup>Whitworth College, 300 W. Hawthorne Road, Spokane, WA, 99251, USA

<sup>c</sup>Intelligent Multimedia Processing Laboratory, School of Computer Science, Telecommunications, and Information Systems, DePaul University, Chicago, Illinois, 60604, USA

## ABSTRACT<sup>†</sup>

Current segmentation techniques require user intervention to fine-tune thresholds and parameters, plot initial contours, refine seed placement, and engage in other optimization strategies. This can cause difficulties for physicians trying to use segmentation tools as they may not have the time or resources to overcome steep learning curves. In order to segment volumetric regions from sequential slices of computed tomography (CT) images with minimal user intervention, we propose an algorithm based on volumetric seeded region growing that employs an adaptive and prioritized expansion. This algorithm requires a user only to identify a voxel in an organ to perform volumetric segmentation. This approach overcomes the need to manually select threshold values for specific organs by analyzing the histogram of voxel similarity to automatically determine a stopping criterion. The homogeneity criterion used for region growth in this approach is calculated from volumetric texture descriptors derived from co-occurrence matrices which consider voxel-pairs in a 3-dimensional neighborhood of a given voxel. Preliminary segmentation results of the kidneys, spleen, and liver were obtained on 3D data extracted from 700 sequential CT images from various studies collected by Northwestern Memorial Hospital. We believe this approach to be a viable segmentation technique that requires significantly less user intervention when compared to other techniques by necessitating only one user intervention, namely the selection of a single seed point.

**Keywords:** semi-automatic segmentation, volumetric segmentation, region growing

## 1. INTRODUCTION

Image segmentation is by definition a separation of a region of interest from the rest of the data within an image. High-quality segmentation of CT images, particularly in the volumetric sense, is exceptionally appealing to a number of medical imaging and other applications, such as content-based image retrieval, face and hand recognition, object tracking, various forms of 3-dimensional computer-aided diagnoses, improving classification approaches, as well as many others. Specifically within the medical field, while the typical users' expertise may not be computing, there is still a strong requirement for a very high degree of accuracy<sup>8</sup>.

However, many current segmentation algorithms require an excess of user intervention in order to properly segment an organ, and often do not provide the degree of accuracy required within the medical domain. For example, active contours and deformable models have difficulty segmenting regions of interest with intricate boundaries, grey-level threshold techniques can not capture regions of varying textures, and edge finding algorithms can not segment a single region of interest within an image. These various techniques also encumber those within the medical domain, whose expertise is rightly focused in the medical field, as they require user intervention to fine tune thresholds and parameters, plot initial contours, redefine seed placement, and engage in other optimization strategies which require region-dependent *a priori* knowledge<sup>9</sup>. The current unsupervised segmentation algorithms do not consistently uphold the same

---

<sup>†</sup> This material is based upon work supported by the National Science Foundation under Grant No. 0453456.

degree of accuracy as the more cumbersome supervised segmentation algorithms. However, In order to improve workflow and efficiency in the medical domain, physicians must not be required to overcome steep learning curves or with fine-tuning parameters and thresholds on an organ-by-organ basis to receive accurate results.

An important aspect of the medical imaging domain is the existence of multiple 3-dimensional image modalities. Along with the increased data in 3-dimensional imaging studies comes a need for complex volumetric segmentation. The results from volumetric segmentation can not only be used for new intentions such as 3-dimensional visualization, but can also be an improved preprocessing step for other classification and analysis algorithms, as it may provide more accurate results which take advantage of the additional data.

According to Smith and Chang<sup>11</sup>, texture refers to a visual pattern that demonstrates properties of homogeneity. Previous studies have demonstrated the ability to perform pixel-level texture extraction within 2-dimensional CT studies<sup>5</sup>. Now that CT studies are isotropic in all three dimensions, it is sought-after to extend this work into the third dimension for additional texture extraction. In order to accomplish this, we implemented a voxel-level texture extraction using Haralick's<sup>3</sup> co-occurrence model on a volumetric co-occurrence<sup>6</sup> matrix extracted from a 3-dimensional neighborhood surrounding each voxel. From these texture features we then produced a homogeneity criterion, which is used to differentiate a region of interest.

The approach used to segment each region of interest is a modified form of seeded region growing, the original being proposed by Adams and Bischof<sup>1</sup>. The modified form used here involves an adaptive and prioritized search as well as voxel-level texture feature extraction to determine the homogeneity criterion. The texture features are dynamically extracted from each voxel as they are included in the region and are used to determine the texture features which characterize the region of interest as it grows. With this, it is possible to keep a prioritized list of all pixels adjacent to the region of interest, ordered by their similarity to the texture features of the region average<sup>2</sup>. By continuously adding the voxel from the front of this list, it is ensured that the most similar voxel is always included first, while preserving the integrity of the region average. Also from this list we are capable of producing a histogram of all voxels included in the region and adjacent to it. From this we can visualize the frequency of the similarities between visited voxels and the region average. We will therefore note a bell-shaped formation, representing the given region of similarity, which is the region of interest. The cut-off threshold can then be determined based on this bell-curve, which will be unique to each region, and the most similar voxels will be included in the region until this threshold is met.

It is important to provide a more efficient means of segmentation, which requires minimal user intervention. Because this algorithm is capable of dynamically determining texture features that quantify the user's selected region of interest, and calculate a stopping criterion based on the texture features, the user is required to do nothing more than select the seed voxel. This limited user intervention should be ideal for the medical domain, assuming the results meet the required level of precision.

## 2. METHODOLOGY

### 2.1 Data Compression

The CT data used consisted of 3D images of consecutive 2D DICOM (Digital Imaging and Communications in Medicine) slices, where each individual slice was 512 by 512 pixels with a 12-bit grey-level resolution. Because co-occurrence matrices were to be used in the texture feature extraction, the number of grey-levels was reduced by a factor of 16, leading to 256 intensities, rather than 4096. This greatly reduced computational complexity, as the number of grey-levels determines the size of each co-occurrence matrix. As proposed in a previous image compression study, Clipped Binning can be used to reduce 4096 grey-levels into 256 grey-levels, while still providing adequate discriminatory power between abdominal organs when considering soft-tissues by means of texture features<sup>10</sup>.

Clipped binning linearly reduces a set range of grey-levels to a smaller range, while clipping all values above and below into single bins surrounding the selected range<sup>7</sup>, as shown in Figure 1. Radiologists from Northwestern Memorial Hospital asserted that soft-tissues in CT images were comprised of the grey-level range [856,1368], which was affirmed by Lerman<sup>7</sup> via the use of a k-means clustering algorithm. Clipped binning was therefore used to reduce all grey-levels

above and below the soft-tissue range into the two appropriate bins, and linearly distribute all soft-tissue pixels into the remaining 254 bins. As a result, with the use of the clipped binning technique on our 4096 grey-level DICOM images, grey-levels [0-855] were mapped to grey-level 0, grey-levels [856,1368] were linearly distributed to grey-levels [1,254], and grey-levels [1369,4095] were mapped to grey-level 255.




Figure 1: Clipped Binning

## 2.2 Voxel-level texture extraction

Any region growing algorithm calls for a criterion of homogeneity in order to facilitate the evaluation of voxels to determine if they should be considered part of a region<sup>1</sup>. Voxel-level feature extraction was used to calculate the 3-dimensional texture features for each individual voxel, in order to compare it to the texture features of the region average. This method is similar to the 2-dimensional pixel-level texture extraction as proposed by Kalinin<sup>5</sup>, as it applies to an additional dimension.

In order to calculate the texture features of a given voxel, we consider a 3x3x3 neighborhood surrounding the voxel, from which we produce a single 2-dimensional co-occurrence matrix. This co-occurrence matrix will be of size n by n, where n is the number of grey-levels in the CT images. Our volumetric co-occurrence matrix sums intensity value pairs, such that the location [x,y] will be the accumulation of voxel pairs of intensities x and y. In determining the voxel pairs for our co-occurrence matrix, we orderly iterate through the 3x3x3 neighborhood such that there is no repetition in the voxel pair calculation.

From this co-occurrence matrix, we calculate four texture descriptors as proposed by Haralick<sup>3</sup>: entropy, mean, variance, and cluster tendency. The definitions of these texture descriptors can be seen in Table 1. Because the texture features of an individual voxel were to be compared with the texture features of the region average, it was necessary to determine the distribution of each texture feature throughout our study. To determine the distribution of each texture feature, we recorded approximately 10,000 random voxels as samples, and calculated their texture features. Within our study, we found the distributions of variance and mean to characterize a fairly normal distribution, while cluster tendency and variance portrayed exponential distributions. In order to normalize the distributions of cluster tendency

Texture Feature	Formula	Description
Entropy	$-\sum_i^M \sum_j^N P[i, j] \log P[i, j]$	Provides a measurement of the randomness of a gray-level distribution
Mean	$\frac{1}{2} \sum_i^M \sum_j^N (iP[i, j] + jP[i, j])$	Provides the mean of the gray-levels in the region
Variance	$\frac{1}{2} \sum_i^M \sum_j^N ((i - \mu)^2 P[i, j] + (j - \mu)^2 P[i, j])$	Provides the variance of the distribution of gray-levels in the region
Cluster Tendency	$\sum_i^M \sum_j^N (i - \mu_i + j - \mu_j)^2 P[i, j]$	Provides a measurement of groupings of voxels with similar gray-levels

Table 1: Texture features

and variance, we took the logarithmic value of each. In this we were capable of determining the four texture features of an independent voxel as it compares with the texture features of the region average.

Once the distributions of the texture features were determined, we were capable of finding a quantitative measurement to find the probability between two values of texture features. As such, we calculated the Mahalanobis distance of each texture feature of an individual voxel from the corresponding texture feature of the region average. With the four Mahalanobis distance values, we computed the Manhattan distance, to achieve the quantitative measurement of the similarity of an individual voxel and the region average, as shown in Figure 2.




Figure 2: Mahalanobis distance

### 2.3 Region Growth

Once there is a means of determining the similarity between a voxel and the region average, we can begin the region growth. Following the user's selection of a single voxel as the seed, the region average will be defined as the seed's texture alone. The algorithm then proceeds to evaluate voxels adjacent to the region, compute their similarity to the region, and maintain a prioritized list of these voxels<sup>2</sup>. It then selects the most similar voxel from the list, includes it in the region, recalculates the region average, and includes the voxel's neighbors in the prioritized list. This process is then repeated, continuously spreading to the most similar voxels first.

After the user has selected a single voxel as the seed, it is necessary to allow an initial period of region growth with no stopping criterion. While this initial growth is much smaller than any region of interest, it will facilitate the development of the region average. This is very beneficial, as the region average will better reflect the texture of the selected region, rather than the individual voxel selected<sup>12</sup>. After the initial growth, the prioritized list of voxels, as well as voxels already in the region, can be used to generate a histogram which displays voxels as sorted by their similarity to the region average. By evaluating such a histogram at various points in time, we are able to determine the characteristics of region growth, identify region boundaries, and formulate a stopping criterion. An example of such a histogram can be seen in Figure 3. After the initial period of growth, this histogram will begin representing a bell-

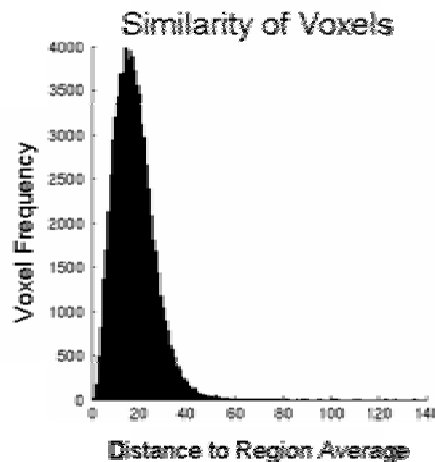


Figure 3: Histogram of a bell-shaped curve denoting a region of interest

shaped curve. At any given time, this bell-shaped curve effectively characterizes a region of interest, with the right-most voxels not yet added to the region as a result of their dissimilarity. The right half of the bell-curve illustrates a decrease in voxel occurrences as the dissimilarity of the voxels from the region average increases. This development from a large number of voxels similar to the region average to a lower number of dissimilar voxels can be interpreted as the border of a region of interest.

Seeing as the right tail of the histogram denotes the border of a region, it would appear that a stopping criterion could simply be based on growing right up to this boundary. However, continuing region growth completely up to the apparent borders continues development of the region average, and its histogram representation, and can lead to leakage. We claim, however, that it is possible to segment a region before filling every voxel up to the border. We claim that you can instead find the core texture of the region, then apply morphological operators to fill in holes and soften edges<sup>4</sup>. With this in mind, we define a threshold value relative to the right tail of the bell-curve, such that a voxel with similarity less than the value will not be included in the region.

Because the initial periods of growth may not be enough to characterize a full region, it is necessary for us to dynamically redefine our threshold value. The threshold value is first defined after the initial growth period as the first bin which is 90% less than, and to the right of, the peak. Seeing as the earlier stages of growth do not yet characterize the region, the threshold value is again calculated when the most similar voxel is more dissimilar than the current threshold. Once recalculated, the threshold value will portray all region growth up to that point. If the most similar voxel still exceeds the newly calculated threshold, then we have truly met our stopping criterion, and the region growing process is concluded. The majority and shape of the region has been determined, and morphological operators are applied. Dilation is used to fill in dense portions of the segmented region, and blurring rounds the edges, resulting in a fully segmented region.

Figure 4 illustrates a portion of the user-interface developed for the algorithm. Note the two vertical bars within the histogram. The right-most bar portrays the current threshold, and the left-most bar portrays the most recently evaluated voxel. The left-most bar will obviously move as voxels are evaluated, and will eventually reach the threshold bar, at which time the threshold value will be recalculated. Also note from this figure that the vast majority of voxels to the right of the threshold bar have not been added to the region, but are instead neighboring voxels waiting in the prioritized list.




Figure 4: An example of a histogram derived from a prioritized list

It's important to note that the threshold value's precision directly depends on the number of bins within our histogram. If too many bins are used in relation to the amount of data, then the bell-curve is very jagged and determining the threshold can be difficult. If too few bins are used, a great amount of precision is lost when determining the threshold. Seeing as the amount of statistical data available for the histogram will increase with the size of the region, it is beneficial to increase the number of bins in the histogram, to benefit the growing need for, and availability of, precision.

$$numOfBins = 75 + 10 * \left[ \frac{\log\left(\frac{numOfVoxels}{1000}\right)}{\log(2)} \right]$$

Equation 1: Equation for the logarithmic bin inflation

Equation 1 portrays the equation used for the increased number of histogram bins. An initial amount of 75 bins is then increased logarithmically as the region expands. This accompanies greater precision with the growing amount of data. The logarithmic increase of histogram bins will provide increased precision for a growing region, but will not result in an excess of bins for enormous regions such as a liver.

### 3. RESULTS

Our segmentation results were obtained on 3D data extracted from 700 sequential CT images from various studies collected by Northwestern Memorial Hospital. The 3D image data consisted of consecutive 2D DICOM (Digital Imaging and Communications in Medicine) slices, each slice being of size 512 by 512 with 12-bit gray level resolution. The following figures were segmented by our algorithm via the selection of a single voxel and no other user intervention or modification of parameters. This only involves the user scrolling through the list of CT images and clicking within a region of interest, as shown in the user interface of Figure 5. Each of the following figures is divided into three different sections. Moving from left to right, those sections are the original CT image, the image with a boundary around the segmented region computed by our algorithm, and the isolated region segmented by our algorithm. Although these regions were segmented volumetrically, only a single slice of the segmentation is shown.




Figure 5: The user interface implemented in the algorithm

Figure 6 depicts one slice of the 3-dimensional results of a segmented left kidney. The two holes in the segmentation are representative of the renal pelvis which is not the goal of our segmentation. This example shows that our algorithm is capable of successfully segmenting organs that have not only external boundaries but internal boundaries as well. It should be noted that these kinds of phenomena prevent active contours and deformable models from proper segmentations.



Figure 6: One slice of the 3-dimensional results of a segmented left kidney

Figure 7 depicts one slice of the 3-dimensional results of a segmented right kidney. It is noteworthy that our segmentation algorithm did not spread from the kidney to the renal vein, which is as intended. This example also demonstrates the ability to capture irregular shapes.




Figure 7: One slice of the 3-dimensional results of a segmented right kidney

Figure 8 depicts one slice of the 3-dimensional results of a segmented spleen.




Figure 8: One slice of the 3-dimensional results of a segmented spleen

Figure 9 depicts one slice of the 3-dimensional results of a segmented liver. In this figure it is important to note that the portal vein, seen in the middle of the liver, was accurately classified as not part of the region.




Figure 9: One slice of the 3-dimensional results of a segmented liver

A large advantage of volumetric segmentation is how well it lends itself to 3-dimensional visualizations. Figure 10 is an example of a kidney segmented by our algorithm and depicted within 3D Slicer, a 3D visualization program. This can be of particular use in the medical domain, as 3-dimensional examination of individual abdominal organs can be obtained by merely clicking the desired region, and using a standard 3-dimensional visualization application on the segmentation results.



Figure 10: A 3-dimensional rendering of a segmented kidney

Figures 11 and 12 are visualizations of a set of segmented kidneys and a spleen within 3D Slicer. The segmentation results of our algorithm were used as coordinates in 'highlighting' the selected organs. This demonstrates yet another advantage and use of volumetric segmentation in the medical domain.






Figure 11: A 3-dimensional visualization of a set of segmented kidneys and spleen




Figure 12: A 3-dimensional visualization of a set of segmented kidneys and spleen

## 4. CONCLUSIONS

The immediately observable benefit of our algorithm is the production of quality volumetric segmentation with a parameter-less user-interface, which requires only a selection of a single seed voxel as user intervention. In particular, the nearly automatic nature of our algorithm allows novice computer users to generate quality segmentation results. This is specifically appealing for the medical domain, which greatly benefits from decreased user intervention while still providing quality results. The results can then be used in a variety of medical diagnoses, or other areas such as generating databases of organs for later similarity-based retrieval.

However, this approach does come with drawbacks. Due to the massive amount of data in isotropic or nearly isotropic volumetric studies, segmentation of very large organs, such as a liver, can take hours. Although speed can be optimized in a development environment that is not research-based, there will still be the computational costs of keeping a prioritized list of all neighboring voxels, and other expensive computations. It is also important to note that the algorithm assumes a criterion of texture homogeneity for a region. Therefore, it will be difficult to segment organs of more complex texture, such as the heart.

For future work it would be worthwhile to experiment with new texture descriptors, or new combinations of texture descriptors, as some may be more discriminating for specific organs. Alternate image compression strategies such as continuously clipped binning<sup>7</sup> may also provide more distinguished grey-levels between regions of interest, or help to lower computational complexity by reducing the number of grey-levels further. We are also investigating the potential for unsupervised segmentation which would make use of an autonomous multi-seed placement strategy.

## References

1. R. Adams, L. Bischof, "Seeded Region Growing", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16 no. 6, pp. 641-647, 1994.
2. M.F. Cohen, J. Painter, M. Mehta, K. Ma, "Volume Seedlings", *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, pp. 139-145, 1992.
3. R. M. Haralick, K. Shanmugam, I. Dinstein, "Textural features for image Classification", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. Smc-3 no. 6, pp. 610-621, 1973.
4. R. Huang, K. Ma, "RGVIs: Region growing based techniques for volume visualizations", *Proceedings of the 11<sup>th</sup> Pacific Conference on Computer Graphics and Applications*, pp. 355-363, 2003.
5. M. Kalinin, D. S. Raicu, J. D. Furst, D. S. Channin, "A classification approach for anatomical regions segmentation", *The IEEE International Conference on Image Processing (ICIP)*, 2005.
6. A. S. Kurani, D. H. Xu, J. D. Furst, D. S. Raicu, "Co-occurrence matrices for volumetric data", *The 7<sup>th</sup> IASTED International Conference on Computer Graphics and Imaging – CGIM 2004*, 2004.
7. R. Lerman, D. S. Raicu, J. D. Furst, "Contrast enhancement of soft tissues in computer tomography images", *Proceedings of SPIE Medical Imaging 2006*.
8. T. S. Newman, N. Tang, S. Bacharach, P. Choyke, "A volumetric segmentation technique for diagnosis and surgical planning in lower torso CT images", *Proceedings of the 1996 International Conference on Pattern Recognition*, pp. 553-557, 1996.
9. S. D. Olabariaga, A. W. M. Smeulders, "Interaction in the segmentation of medical images: A survey", *Medical Image Analysis*, vol. 5, pp. 127-142, 2001.
10. D. S. Raicu, S. Handrick, B. Naimipour, J. D. Furst, "Binning strategies evaluation for tissue classification in computer tomography images", *Proceedings of SPIE Medical Imaging 2006*.
11. J. Smith, S. Chang, "Automated binary texture feature sets for image retrieval", *IEEE International Conference on Acoustics*, vol. 4, pp. 2239-2242, 1996.
12. Yongming Li, Dongming Lu, Xiqun Lu, Jianming Liu. "Interactive Color Image Segmentation by Region Growing Combined with Image Enhancement Based on Bezier Model," *icig*, pp. 96-99, Third International Conference on Image and Graphics (ICIG'04), 2004.