Determination of Intensity Thresholds via Shape Gradients

Roger Tam and Alain Fournier*

Department of Computer Science,
University of British Columbia

Abstract

Shape analysis is a vital aspect of medical imaging. The shapes of objects in an image provide high-level information that is essential for many image processing tasks. Accurate analysis of medical images is often dependent upon an appropriate greyscale thresholding of the image for reliable feature extraction. The determination of object thresholds can be a time-consuming task because the thresholds can vary greatly depending upon the quality and type of image. Thus, an efficient method for determining suitable thresholds is highly desirable. This paper presents a method that uses shape information to accurately determine the intensity ranges of objects present in a greyscale image. The technique introduced is based on the computation of the *shape gradient*, a numerical value for the difference in shape. In this case, the difference in shape is caused by the change in threshold value applied to the image. The use of this gradient allows us to determine significant shape change *events* in the evolution of object forms as the threshold varies. The gradient is computed using *Union of Circles* matching, a method previously shown to be effective in computing shape differences. We show the results of applying this method to artificially computed images and to real medical images. The quality of these results shows that the method is potentially viable in practical applications.

CR Categories: I.4.6 [Image Processing and Computer Vision]: Segmentation—Edge and feature detection, Pixel classification I.4.7 [Image Processing and Computer Vision]: Feature Measurement—Feature representation, Size and shape I.4.10 [Image Processing and Computer Vision]: Image Representation—Morphological, Volumetric

Keywords: feature extraction, feature matching, image processing, medical imaging, segmentation, shape representation, shape measurement, thresholding, Union of Circles

1 Introduction

An important task in medical imaging is to accurately assign percentages of different types of tissue (e.g., bone, fat, soft tissue, etc.) to each pixel in an image [15]. For most applications, this material classification process is a vital step required for further processing of the image, such as feature extraction. For example, the quality of volume rendered images depends strongly on the accuracy of the transfer functions used [5]. The determination of intensity thresholds is often an important part of material classification. The most common class of methods used for determining thresholds in medical imaging is histogram-based analysis. This form of analysis can be problematic because histograms often do not reveal obvious differences between the various materials present in an image, especially if the image quality is low (e.g., noisy or blurry). In addition, the use of different windowing parameters [14] can cause images of the same anatomical structure to have very different histograms.

A major reason histogram-based techniques often do not work well is that histograms do not capture any spatial characteristics of the objects represented. The method introduced here utilizes shape information along with intensity data to precisely identify thresholds of interest. The core idea of our method is to measure the differences in the shapes of objects as the threshold is varied. (Note that an "object" in this case can be comprised of two or more spatially distinct components of the same tissue type. For example, two kidneys in an image would be considered collectively as one object.) The primary contribution of this paper is the use of the *shape gradient*, the amount of shape difference caused by a given change in threshold value, to determine the occurrence of significant shape change *events* in the given intensity range. These events determine the thresholds that we should use for material classification. More precise definitions for shape gradient and shape event are found in Section 2.

The technique we use for measuring shape differences across intensity levels is based on *Union of Circles* (*UoC*) matching, a method shown by Ranjan and Fournier [20] to be effective for shape comparison. Ranjan and Fournier showed the UoC representation to be robust, stable and efficient. We show that our method effectively detects the presence of significant objects ("significant" in this case means a human would recognize them as objects) and determines their upper and lower thresholds with no *a priori* assumptions about the objects represented.

Figure 1 provides a motivating example. The image of the sharp 'E' on the left has only two intensities: 25 inside the 'E' and 0 outside (the image shown here has been brightened for printing). Thus, the intensity histogram is very simple and finding thresholds for this image is trivial. In contrast, Figures 1b and 1c show a blurred and noisy version of the 'E', respectively. Figure 2 shows the histograms of these two images. Even though the images show essentially the same shape, their histograms are very different, and it is not at all obvious what thresholds would best define the object's intensity range.

Figure 3 shows the graphs of how the shape gradient varies with intensity for the blurred and noisy 'E' images. As explained in Section 2, a minimum in a shape gradient plot marks the lower threshold of a significant object. The plots in this case have obvious minima at 14 for the blurred 'E' and at 16 for the noisy 'E'. These minima indicate good thresholds to use based on the shape of the object. Figure 4 shows the images in Figure 1 with their respective thresholds applied. Clearly, the method has found thresholds suitable for bringing out the shape of the object while minimizing the artifacts caused by blurring and noise. However, these are only simple test cases designed to explain the motivation for this work. We show in Section 3 that our method also works well for actual medical images.

1.1 Related Work

Intensity thresholding is considered to be a low-level segmentation technique, and is at one end of a broad spectrum of methods for pixel classification. (For a recent survey of biomedical image segmentation techniques, refer to Acharya and Menon [1].) Low-level techniques make very few assumptions about the input data and are very generally applicable.

Although our algorithm is presented in this paper as a method for determining intensity thresholds, an objective normally associated

* {rtam|fournier}@cs.ubc.ca

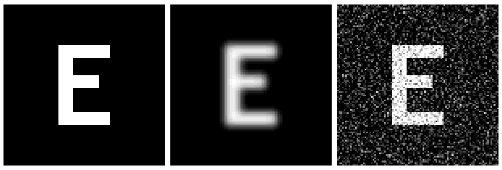

(a) (b) (c)

Figure 1: (a) Sharp 'E' (b) Blurred 'E' (c) Noisy 'E'

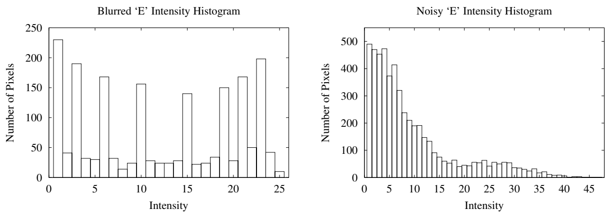

(a) (b)

Figure 2: (a) Intensity histogram for blurred 'E' (b) Intensity histogram for noisy 'E'

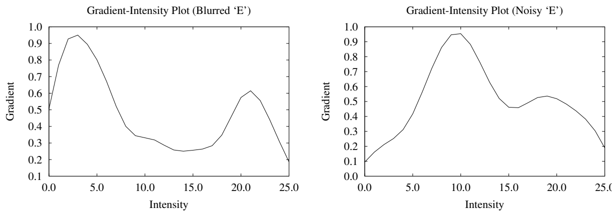

(a) (b)

Figure 3: (a) Shape gradient plot for blurred 'E' (b) Shape gradient plot for noisy 'E'

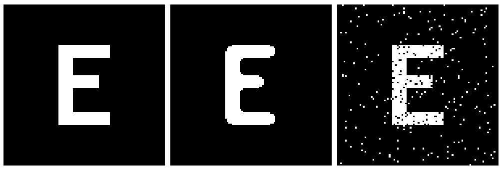

(a) (b) (c)

Figure 4: (a) Sharp 'E' (b) Thresholded blurred 'E' (c) Thresholded noisy 'E'

with low-level techniques, it can actually be best compared with other methods commonly classified as "medium-level" approaches. These methods, like ours, aim to extract and utilize high-level shape information while trying to minimize assumptions about the input data. Most medium-level methods output a set of edges or contours that represent the boundaries of objects in the image. Although our algorithm is capable of outputting object contours, our focus in this paper is the determination of intensity thresholds.

Many medium-level techniques use feature clustering, which first extracts a number of primitives, most commonly edges, then tries to connect these primitives together to obtain models of the objects. Edge-based approaches often have problems with being unstable, and depend strongly on the quality of the edge detector used. Two of the more successful techniques that utilize the edge-connection approach are Wu and Leahy [24] and Mallat and Zhong [10].

Many of the other more successful medium level approaches use *deformable models* [11]; early examples include energy minimizing snakes attracted to image features such as lines and edges (Kass *et al.* [7]) and deformable superquadrics by Terzopoulos and Metaxas [21], which use global shape parameters with local degrees of freedom based on elastic properties and external forces. Most of these techniques have to impose continuity (geometric or topological) or smoothness constraints on the models, which limits the types of objects that can be represented with any given set of parameters. Recent papers by McInerney and Terzopoulos [12, 13] address the issue of topological adaptability of snakes and deformable surfaces. Sethian's *level sets* [9] approach, which uses a deformable model that makes no assumptions about the object's topology, is similar to our technique in that the image is analyzed over a number of intensity levels.

Another group of popular shape models are based on a medial axis representation, pioneered by Blum [2]. The most comprehensive approach to date for medical image processing is by Pizer *et al.* [17], who represent shapes using interconnected figures, where a *figure* is a whole object, the main part of an object or a protrusion or indentation in another figure. Notable features of this representation include a shape metric for the measurement of shape variation, four levels of coarseness for multiscale processing and positional tolerance for enhanced stability. The major drawback of the Pizer model for the application in this paper is that it is much more mathematically complicated than necessary (see [16] for the mathematical details).

Our algorithm is based on the Union of Circles approach to shape representation that was pioneered by Ranjan and Fournier [20, 18, 19]. For use in our technique, UoC models have a number of advantages over the shape models mentioned above:

1. UoC models are known to be stable with respect to changes in the input data. This allows the algorithm to handle a wide variety of images.

2. There is an effective shape measure that can be used to quantify the difference between any two UoC's regardless of their differences in topology or number of components. As we will show in Section 2, this is an important property for obtaining the shape gradient between threshold levels.

3. UoC models are simple and efficient.

The shape extraction and matching performed by our algorithm share many of the same goals as image registration, non-rigid motion tracking and image morphing. For a good overview of registration techniques, the reader is referred to the survey papers by Brown [3] and Maintz and Viergever [8]. The latter paper focusses more on issues specific to the medical image domain. Kambhamettu *et al.* [6] have written an overview paper on non-rigid motion analysis. Wolberg [22, 23] has written extensively on morphing.

2 Methodology

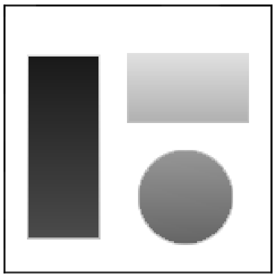

Figure 5: Simple test image

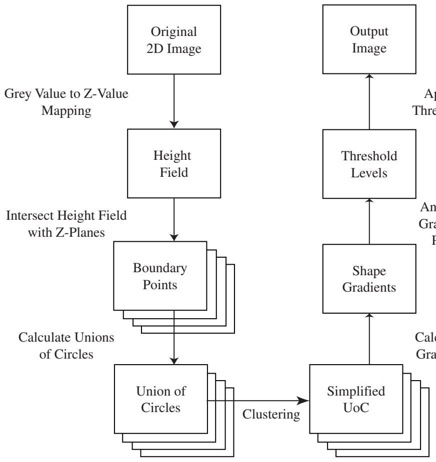

Figure 6: Overview of shape-driven threshold determination algorithm

This section describes the details of our algorithm for detecting significant objects and determining the intensity range of each detected object. The simple image shown in Figure 5 is used as an example to explain the processes involved. This sample image contains three objects, each having a linear intensity gradient.

As illustrated in Figure 6, the main steps of our algorithm are:

1. Generate boundary points.
   (a) Generate a height field from the image's intensity data by mapping the value at each pixel to a z-value.
   (b) Intersect the height field with n planes $z = I_i$, $i = 1 \ldots n$. This results in a set of boundary points for each of n grey levels.
2. Calculate a Union of Circles for each set of boundary points.
3. Simplify each Union of Circles by clustering.
4. Calculate the shape gradient between successive levels, using the Union of Circles method of shape matching.
5. Determine the thresholds of interest from the maxima and minima in the shape gradient data.
6. Apply thresholds to generate an output image, or use thresholds for other processing.

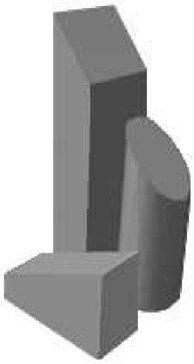

Figure 7: Volume created from the image in Figure 5

2.1 Boundary Point Generation

Creating a height field from a greyscale image is relatively straightforward. All that is required is a mapping from the intensity value at each pixel to a z-value in the height field. A simple linear mapping suffices in most cases, and is what we use for the test cases presented in this paper. The points in the height field are then used as the boundary points of a volume that is bounded below by the image plane. This volume can be comprised of a number of disjoint components. Figure 7 shows the volume created from the example in Figure 5.

The resulting volume is then intersected with n different z-planes. The values used for z are $\{I_i,\ i = 1, \ldots, n : I_{min} <= I_i <= I_{max},\ I_{i+1} = I_i + I_{incr}\}$, where I_{min} and I_{max} are the minimum and maximum intensity values of the image, and I_{incr} is the intensity increment (typically 1) from one level to the next. The intersection of each plane with the volume results in a set of boundary points at that level.

2.2 Union of Circles Generation

The next step in the algorithm is to generate a UoC for each level. There are three basic steps to forming a UoC from a set of boundary points (further details can be found in [19, 20]):

1. Compute the Delaunay triangulation of the point set.
2. Compute the circumscribing circle of each triangle.
3. Discard all circles that are outside of the 2D region defined by the intersection of the plane with the volume.

The remaining circles form the UoC. Figure 8 shows the UoC generated by intersecting the volume in Figure 7 at $z = 10$. Because all three objects in the image have intensities greater than 10, three sets of circles are present at this level.

2.3 Union of Circles Simplification

The simplification process aims to reduce the number of circles as much as possible while preserving the shape of the represented objects. The simplification procedure is an important step in increasing the stability of the UoC representation because it removes redundancies from the model. The result is that similarly shaped objects have similar UoC's, regardless of variations in the number and position of boundary points (such as those caused by discretization or noise).

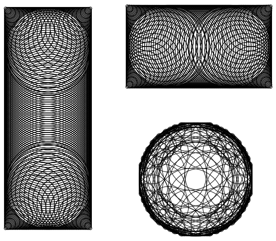

Figure 8: Union of Circles (640 circles) generated by intersecting the volume in Figure 7 with the plane $z = 10$.

Figure 9 shows a simplified version of the UoC in Figure 8. The original UoC has 640 circles, with 150 circles representing the disk in the lower right. Clearly, these shapes can be adequately represented with far fewer circles. The redundancy is primarily caused by the density of the boundary points used for calculating the UoC. In contrast, the simplified UoC only has 104 circles, with a single circle representing the disk shape. The simplified UoC is a more efficient representation and is more useful for comparison with other UoC's.

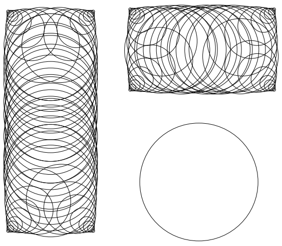

Figure 9: Clustered Union of Circles (104 circles)

The simplification algorithm works by taking clusters of circles within the UoC, and replacing each cluster with a single encompassing circle. Hence the process is called *clustering*. The degree of simplification is controlled by a user-set parameter called *sphericity*, which is a measure of how well a set of circles can be modelled by a single circle. (A more thorough discussion of sphericity and clustering can be found in [19, 20].) We have found that a sphericity of 0.96 works well for all of the images we have tested to date, which include a wide variety of computed tomography (CT) and magnetic resonance imaging (MRI) images.

Figure 10 shows an example of the effectiveness of the clustering method in preserving shape while reducing the number of circles. Using a sphericity of 0.96, the number of circles in the hand is reduced by more than 50%. It is worth noting that small circles are still present where necessary to preserve shape, such as in the fingers. In areas of relatively low detail, such as in the palm, the number of circles is greatly reduced.

2.4 Shape Gradient Computation

The next main step in our algorithm is the computation of shape gradients between UoC's on successive levels. The shape difference between two UoC's is computed by matching circles between the two models using a specially defined distance measure. The shape distance between the two UoC's is taken to be the average of the distances between all matched pairs of circles. Figure 11 shows three UoC's computed from the image in Figure 5 at three different levels.

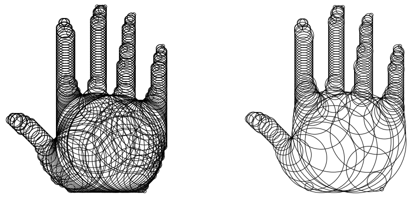

(a) (b)

Figure 10: Clustering using a sphericity of 0.96. (a) Original UoC (373 circles) (b) Clustered UoC (180 circles)

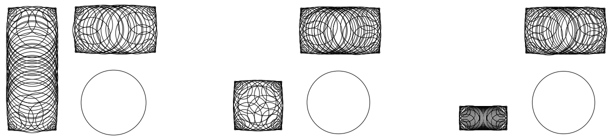

(a) (b) (c)

Figure 11: Union of Circles of Figure 5 at three intensity levels (a) 25 (b) 55 (c) 65

The first step in the matching process is the calculation of the distances $d(a, b)$ between every a and b, where a is a circle in the first UoC, and b is a circle in the second. Given that a has centre (x_a, y_a) and radius r_a and b has centre (x_b, y_b) and radius r_b, the distance function is given by:

$$d(a, b) = w_p d_p(a, b) + w_s d_s(a, b) + w_f d_f(a, b)$$

where $d_p(a, b) = (x_a - x_b)^2 + (y_a - y_b)^2$, $d_s(a, b) = (r_a - r_b)^2$, and $d_f(a, b)$ is the *feature* distance between a and b. The definition of a feature in this case is a mathematical relationship between a circle and its three largest neighbours. (For a precise definition of the feature of a circle, we refer the reader to [20].) Thus, we define the distance between two circles as a function of the differences in their positions, sizes and neighbourhood information.

The weights w_p (position), w_s (size) and w_f (feature) are chosen by the user. For our purposes, because we are interested in the relative rather than absolute values for the shape gradient, we can assign some reasonable value to each weight and use the same three values for all images. We have found that simply setting all three weights to 1.0 works well.

After the distances between all circles in the two levels have been calculated, a weighted bipartite graph is built where the nodes correspond to the circles, and the weights on the edges are the distances between them. A maximum match is computed such that the sum of the distances between all matched pairs is a minimum. The final shape distance is the average distance calculated over all of the matched pairs in the maximum match.

An analysis of how the shape gradient changes as the threshold varies allows the user to determine the ranges of intensity where one or more objects in the image are actively changing shape, as well as the ranges where the shapes are relatively stable. A graph of gradient versus intensity, in which for each threshold value I_i the shape difference between the UoC's at I_i and I_{i+1} is plotted, allows the "thresholds of interest" to be determined easily. We refer to this graph as the *matched gradient* plot.

In most cases, the number of circles does not remain constant across levels. Therefore, the matching process can leave a number of circles unmatched at each level. A circle that exists at one level but "disappears" at the next is considered to have shrunk in place to a radius of zero. A plot of $\bar{r^2}$ versus intensity, where $\bar{r^2}$ is the average of the radii squared of unmatched circles at each level, may contain additional shape change information not captured by the matched gradient plot. This *unmatched gradient* plot is especially useful in cases where one or more significant objects in the image have (close to) completely uniform intensity, as these objects would disappear between UoC's on successive levels without affecting the matches in the other objects. An unmatched gradient plot would have very distinct spikes at such intensities. Thus, even though for most real-world medical images the matched gradient plot alone is sufficient for threshold determination, the unmatched gradient plot is sometimes useful for providing complementary information.

2.5 Shape Gradient Analysis

The following observations are useful when analyzing a matched gradient plot:

- Minima indicate relative stability in the shapes of objects as the threshold is varied.
- Maxima indicate shape change events. A sudden rise in the matched gradient occurs when the threshold reaches a point where a small increase in intensity causes a significant object to breakdown and/or distort.
- From the above, we can conclude:
    1. A minimum or the point at the beginning of a peak marks the lower threshold of a significant object.
    2. The point immediately at the end of a local peak marks the upper threshold of a significant object.
- A wide peak indicates the object of interest spans a relatively large intensity range; a sharp spike means the object spans a narrow intensity range.

To find the intensity ranges of significant objects in an image, we first divide the intensity range of the image into subranges using the minima as boundaries. This creates a number of "buckets", each with an upper and lower threshold. Each bucket contains one significant object. To define the intensity range of each object more precisely, we use the points at the beginning and end of each peak in each bucket, if there are any such points that are not also minima, to further resize the bucket. The precise location of each of these points is defined by finding the local maximum change in slope. The following example illustrates this process.

Figure 12 shows the shape gradient plot computed from the image in Figure 5. The minima at $75 < I < 100$ and $150 < I < 175$, where I is intensity, clearly divide the graph into three buckets ($I \leq 75$, $100 \leq I \leq 150$, $I \geq 175$). The point at the beginning of the first peak ($I = 25$) and the point at the end of the last peak ($I = 225$) are used to resize the buckets to the three ranges ($25 \leq I \leq 75$, $100 \leq I \leq 150$, $175 \leq I \leq 225$) that correspond to the three objects present in the image. Within the first peak, there is a sharp portion that starts at $I = 55$. This can be explained by considering how the long rectangle on the left side of the image changes shape as the threshold increases from $I = 25$ to $I = 75$. As shown in Figure 11, the rectangle, which has a linear gradient along its length, starts to shorten when the threshold gets above $I = 25$, and continues to do so until it disappears at $I = 75$. The sharper portion of the peak in the matched gradient starts when the rectangle becomes a square (Figure 11b), and reaches its maximum as the square becomes a rectangle perpendicular to the original. This transition is detected as a significant shape change event by the UoC matching. Depending on the application, this shape change may be important enough to warrant a further division at $I = 55$.

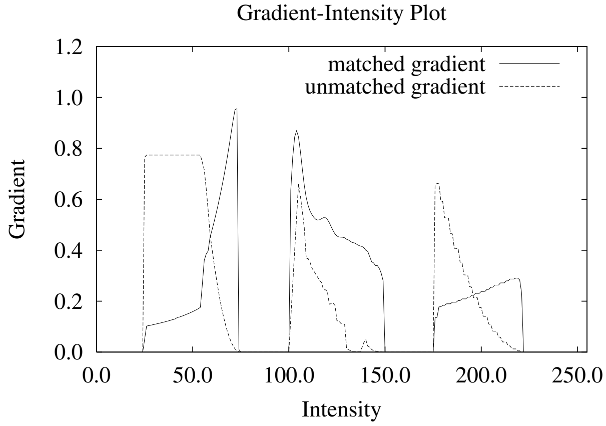

Figure 12: Shape gradient plot computed from the image in Figure 5

In the $100 \leq I \leq 150$ range, there is a local minimum at $I = 115$. However, the amplitude of this minimum is very small relative to the rest of the graph. This brings into question how we should decide which maxima and minima are to be considered "significant". Currently, the user has to manually make these decisions based on the level of detail at which he wishes to analyze the image. Applying a Gaussian filter to the plot to smooth out small maxima and minimum is a very helpful step, but then the choice of the size of the filter kernel becomes important. We are currently working on a method based on the Gaussian pyramid [4] to automatically determine the optimal filter width.

3 Results

This section demonstrates the effectiveness of our algorithm by presenting the results of applying our technique to three representative test cases, all of which are actual medical images. The first test case, shown in Figure 13, is an MRI image of a human brain. This image has 128 grey levels.

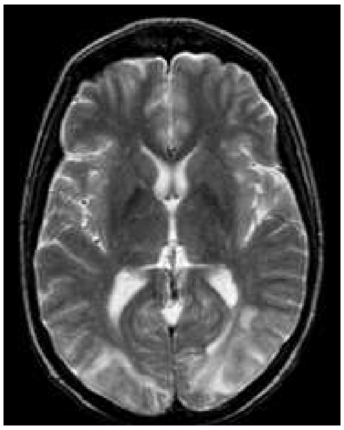

Figure 13: Brain MRI

Figure 14 shows the shape gradient plot computed from the brain MRI. Using the analysis method outlined in Section 2.5, we can clearly identify three significant intensity ranges. The first ($20 \leq I \leq 40$) is the range for the fluid surrounding the brain, and is of limited interest for medical analysis. The other two ranges, labelled R1 and R2 in Figure 14, are for the brain itself and the grey matter inside the brain. Figure 15 shows the two UoC's representing the shapes identified by our algorithm as being stable at $I = 43$ and $I = 100$. (The unmatched gradient confirms stability at those intensities.) Figures 15(a) and 15(b) are clearly good representations of the overall shape of the brain and the grey matter inside.

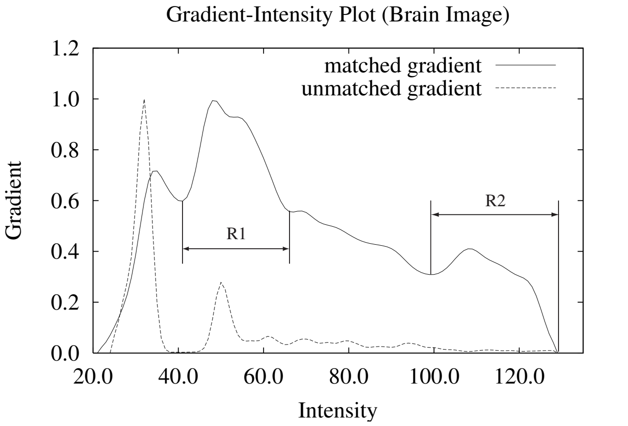

Figure 14: Shape gradient plot computed from the brain MRI image

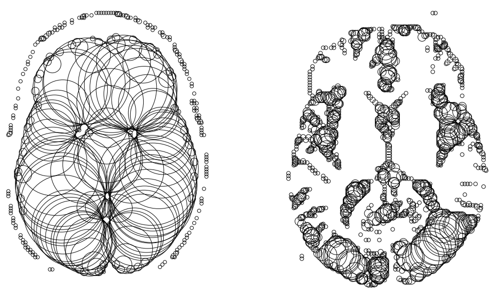

(a) (b)

Figure 15: Unions of Circles computed from the brain MRI at two levels: (a) 43 (whole brain) (b) 100 (grey matter)

The second test case, shown in Figure 16, is a CT image of a person's lower abdomen. The most noticeable structures present are the liver, kidney, small intestine and spine. This image has 128 grey levels.

Figure 17 shows the matched gradient plot from this second test case. Dividing the graph using the minima results in four intensity ranges. The first ($5 \leq I \leq 30$) corresponds to the fluid and soft tissue surrounding the organs. The second ($31 \leq I \leq 51$) corresponds to the liver. The third ($52 \leq I \leq 76$) is associated with the small intestine, and the last ($77 \leq I \leq 109$) is the intensity range for the kidney. Figure 17 shows the CT image with each of the four sets of upper and lower thresholds applied. Again, the algorithm has resulted in effective thresholds for the significant objects in the image. One problem that is apparent with this test case is that the method does not have a distinguishable maximum for the spine. This is because the intensity range of the spine overlaps with those of the liver and small intestine, and since the spine is smaller than the other two structures, it is essentially lost. This kind of problem is common to all algorithms using global thresholds.

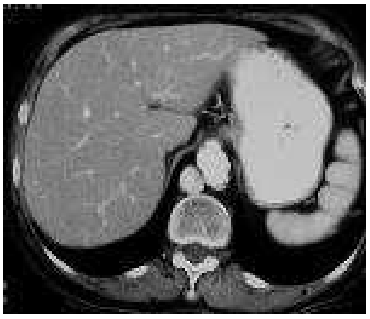

Figure 16: Lower abdomen CT

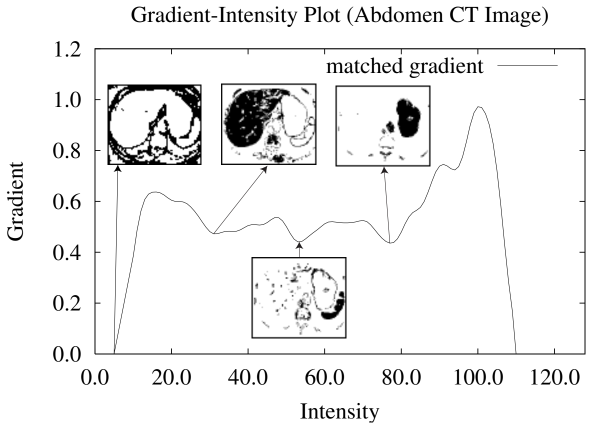

Figure 17: Shape gradient plot computed from the abdomen CT image

The third example, shown in Figure 19(a), is comprised of 10 frames of a time-gated MRI sequence of a beating heart. This set of images is a good example of how the intensity ranges of the same structures can vary greatly across images and make histogram-based analysis very difficult. In this case, these variations are caused by the imaging technique used. Figure 18 shows the shape gradient plot for the first frame. The two minima divide the graph into three segments. The first range ($48 \leq I \leq 153$) corresponds to the walls of the heart, the second ($154 \leq I \leq 212$) to the inner chambers, and the third ($213 \leq I \leq 255$) to blood that is flowing in a direction normal to the image plane and towards the viewer. Table 1 shows the lower thresholds of the three objects of interest. The values for the first threshold ranges from 48 to 51, the values for the second threshold from 120 to 172 and the values for the third threshold from 178 to 226.

Figure 19(b) shows the same heartbeat sequence, colourized according to the extracted thresholds. The heart walls are blue, the inner chambers are light red and the areas of blood flow toward the viewer are dark red. The thresholds clearly define the objects that we wish to visualize. These results demonstrate that our technique is capable of using shape effectively to determine accurate thresholds.

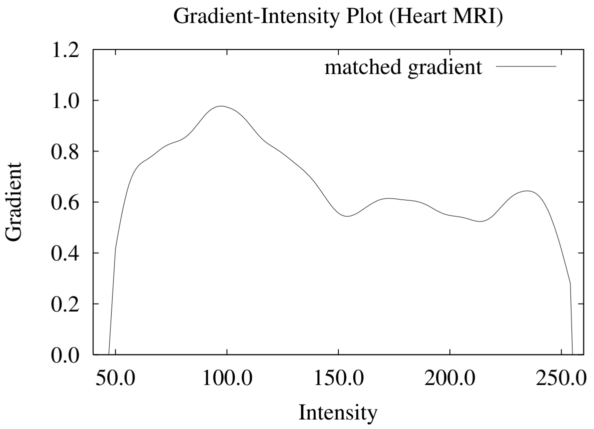

Figure 18: Shape gradient plot computed from the first frame of the heartbeat MRI sequence

Frame	Threshold 1	Threshold 2	Threshold 3
1	48	154	213
2	49	172	201
3	50	140	222
4	50	139	195
5	48	155	195
6	50	146	226
7	51	139	207
8	49	120	178
9	50	125	207
10	50	137	224

Table 1: Thresholds extracted from heart MRI frames

4 Summary and Future Work

We have presented an effective method for detecting and determining the thresholds of significant objects from a greyscale image. We introduced the use of the shape gradient, computed across successive levels of intensity, to identify thresholds of interest. The Union of Circles model is used for the computation of this gradient. The primary advantages of the UoC model over other shape models are stability and the presence of an effective shape measure. We presented the test results of applying our algorithm to artificial images and real medical images to illustrate the advantages and limitations of our technique. These results indicate that the our algorithm is potentially viable in practical applications.

There are a number of directions for future work that we are currently pursuing. A high priority goes to investigating methods for automatically finding the optimal filter size for the gradient curve and for accurately locating points of interest on the curve. Filters other than Gaussian are also being tested.

Although our method gives visually reasonable results, its accuracy needs to be more formally verified before it is clinically applicable. Validation via comparison with thresholds selected by experts (such as radiologists) should done.

Currently, the unmatched gradient is of limited use. Its role is mainly to confirm or reject the results obtained from the matched gradient. Finding a shape measure that combines the matched and unmatched gradients would be preferable.

We are also exploring other ways in which the shape gradient plot can be exploited, such as comparing plots of images of the same object obtained by different scanning modalities, with the objective of determining the optimal thresholds for matching the UoC's extracted from the images. The number of significant shape events present in a plot is typically much lower than the number of pixels or detected edges; this greatly reduces the number of primitives that need to be matched. Therefore, UoC matching should be an efficient method for comparing shapes in different images, or for locating a pre-defined shape in an image.

An extension of the method to 3D using Unions of Spheres [18] to determine appropriate thresholds for extracting isosurfaces would be a straightforward but very useful exercise.

References

[1] R. Acharya and R.P. Menon. A review of biomedical image segmentation techniques. In A. Singh, D. Goldgof, and D. Terzopoulos, editors, *Deformable Models in Medical Image Analysis*, pages 140–161. IEEE Computer Society Press, 1998.

[2] H. Blum. A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, editor, *Models for the Perception of Speech and Visual Form*. MIT Press, Cambridge, MA, 1967.

[3] L.G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.

[4] P. Burt and E. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4):532–540, 1983.

[5] R. A. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. *Computer Graphics*, 22(4):65–74, 1988.

[6] C. Kambhamettu, D. Goldgof, D. Terzopoulos, and T.S. Huang. Nonrigid motion analysis. In A. Singh, D. Goldgof, and D. Terzopoulos, editors, *Deformable Models in Medical Image Analysis*, pages 270–284. IEEE Computer Society Press, 1998.

[7] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: active contour models. *International Journal of Computer Vision*, pages 321–331, 1987.

[8] J.B.A. Maintz and M. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.

[9] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modelling with front propagation: a level set approach. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995.

[10] S. Mallat and S. Zhong. Characterization of signals from multiscale edges. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(7):710–732, 1992.

[11] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.

[12] T. McInerney and D. Terzopoulos. Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Trans. Medical Imaging*, 18(10):840–850, 1999.

[13] T. McInerney and D. Terzopoulos. T-Snakes: Topology adaptive snakes. *Medical Image Analysis*, to appear.

[14] A. Moss and G. Gamsu. *Computed tomography of the body with magnetic resonance imaging*. Saunders Publishing, 1992.

[15] D.R. Ney, E.K. Fishman, and D. Magid. Volume rendering of computed tomography data: Principles and techniques. *IEEE Computer Graphics & Applications*, 10(2):24–32, 1990.

[16] S.M. Pizer, D. Eberly, D.S. Fritsch, and B.S. Morse. Zoom-invariant vision of figural shape: The mathematics of cores. *Computer Vision and Image Understanding*, 69(1):55–71, 1998.

[17] S.M. Pizer, D.S. Fritsch, P.A. Yushkevich, V.E. Johnson, and E.L. Chaney. Segmentation, registration and measurement of shape variation via image object shape. *IEEE Trans. Medical Imaging*, 18(10):851–865, 1999.

[18] V. Ranjan. *A union of spheres representation for 3D objects*. PhD thesis, Department of Computer Science, University of British Columbia, 1996.

[19] V. Ranjan and A. Fournier. Volume models for volumetric data. *IEEE Computer, Special Issue on Volume Visualization*, 27(7):28–36, 1994.

[20] V. Ranjan and A. Fournier. Matching and interpolation of shapes using unions of circles. *Computer Graphics Forum (Eurographics '96 Proc.)*, 15(3):35–42, 1996.

[21] D. Terzopoulos and D. Metaxas. Dynamic 3D models with local and global deformations: deformable superquadrics. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(7):703–714, 1991.

[22] G. Wolberg. *Digital Image Warping*. IEEE Computer Society Press, 1990.

[23] G. Wolberg. Image morphing: a survey. *The Visual Computer*, 14(8/9):360–372, 1998.

[24] Z. Wu and R. Leahy. An optimal graph theoretic approach to data clustering: theory and its application to image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(7):1101–1113, 1993.

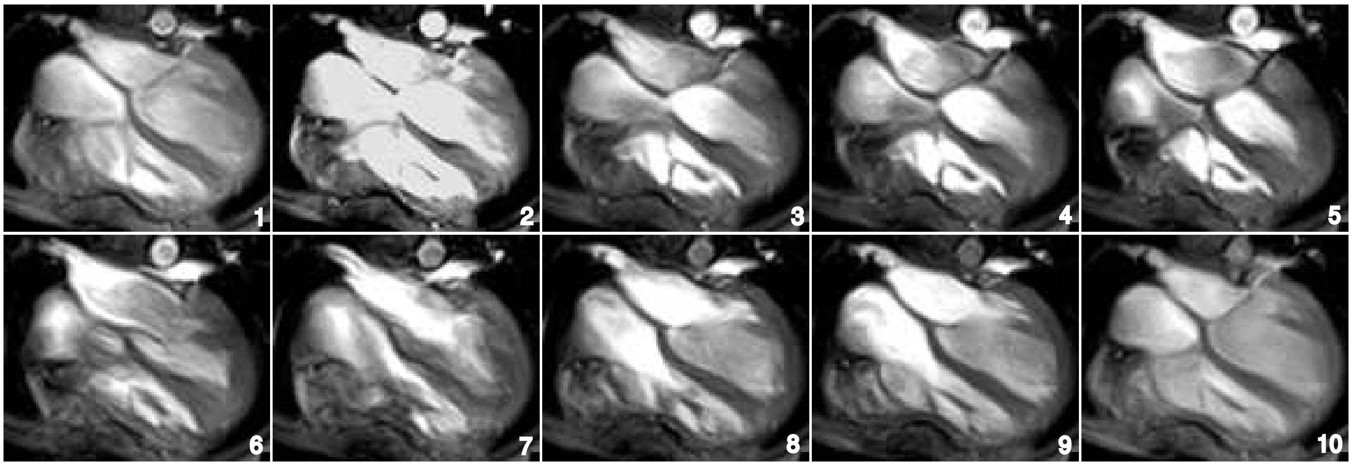

(a)

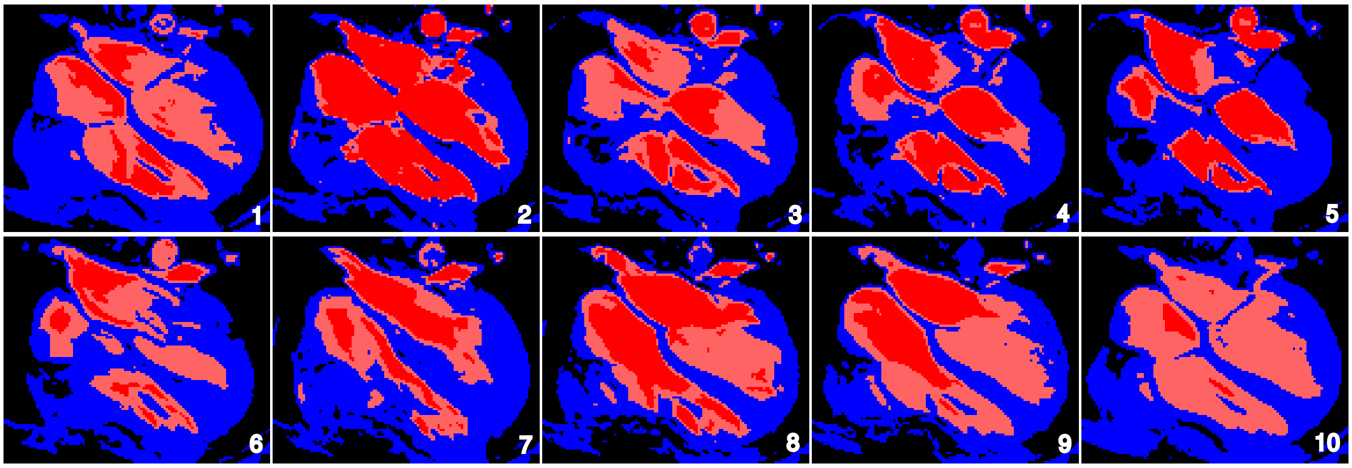

(b)

Figure 19: (a) Time-gated MRI sequence of a heart during a single cardiac cycle (b) Colourized version of the same sequence using shape gradient for pixel classification