Shape Model and Threshold Extraction via Shape Gradients

Roger Tam and Alain Fournier

Department of Computer Science, University of British Columbia
Email: rtam@cs.ubc.ca

Abstract

Shape information is utilized by numerous applications in computer vision, scientific visualization and computer graphics. This paper presents a novel algorithm for exploring and extracting 2D shape information from greyscale images with no *a priori* information about the objects represented. From an image, our algorithm outputs the intensity range spanned by each significant object as well as a versatile shape model of the object that is directly useful for many applications.

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. Our algorithm uses the *Union of Circles* shape representation, which is flexible, stable and has an effective shape metric. The extraction method is consequently robust and works in the presence of noise and other artifacts. We show the results of applying this method to artificially created images and to real medical images.

1 Background and Motivation

The traditional approach to object modelling for shape-related applications such as segmentation and registration focuses narrowly on the goals of the application at hand. As a result, a model developed for a particular application may not capture many shape properties that are useful for other applications. An alternative approach, as exemplified by this paper, aims to develop a more general and flexible shape model that can be used for a broad range of applications. The primary motivation for this approach is to enhance the efficiency of visualization pipelines consisting of multiple shape-related applications.

A general shape model must be able to handle a wide variety of input data. In order to do so, the model must be stable and flexible. Because the core process of many shape-driven applications is shape matching, these properties are of particular importance when considering the shape metrics[1] being applied to the model.

In this paper, we present a novel, efficient and largely automatic method for exploring shape information in greyscale images. Our algorithm effectively extracts from image data significant thresholds and 2D shape models possessing the qualities mentioned above. This method uses the *Union of Circles* (*UoC*) shape representation, which has been shown by Ranjan and Fournier [14] to be stable, flexible and effective for shape comparison. Our extraction algorithm is designed to take advantage of these properties. The main idea of our method is to measure the differences in the shapes of objects in an image as we vary the intensity threshold applied to the image. 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 extracting shape models that represent objects of potential interest to the user. More precise definitions for shape gradient and shape event are found in Section 3.

Figure 1 provides a motivating example that illustrates how the shape gradient method compares to the most commonly used method for determining thresholds, histogram analysis (*e.g.*, [9]). The 'E' on the left has an intensity of 25, with the background having an intensity of 0.[2] Because the boundary of the 'E' is sharp, the intensity histogram

[1] As is frequently done in shape literature, the term *metric* is being used loosely; *pseudo-metric* or *measure* would be more mathematically correct, but we are using *metric* for the sake of consistency.

[2] The pixel intensities of the three images in Figure 1 have been scaled and inverted to increase visibility for the reader.

for this image is very simple and finding thresholds for this object is trivial. In contrast, Figures 1b and 1c show a blurred and noisy version of the ‘E’, respectively.

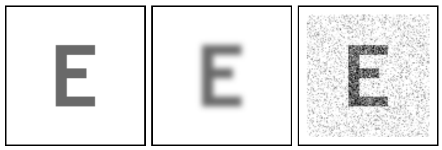

(a) (b) (c)

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

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, especially if the blurring and noise characteristics are not known.

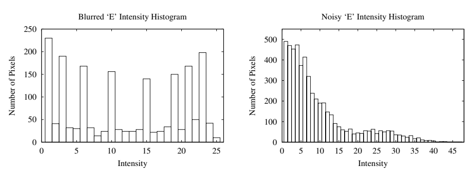

(a) (b)

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

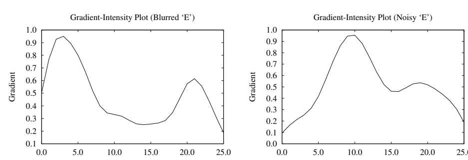

(a) (b)

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

Figure 3 shows the graphs of how the shape gradient varies with intensity for the blurred and noisy ‘E’ images. As explained in Section 3, 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 lower thresholds to use based on the shape of the object.

From the information in the gradient plots, we can compute models of the objects that can be used for subsequent shape-related operations. This simple example only gives a glimpse of the effectiveness of our technique. In Section 4 we show that our method also works well for actual medical images.

2 Related Work

The work presented here is related to many interesting papers dealing with shape modelling for feature extraction. We briefly discuss a representative sample here.

At the low end of the complexity spectrum are feature clustering techniques that do not use a higher-level shape model. The idea is to first extract from the image some simple primitives, most commonly edges, then connect them together to form representations of the objects in the image. Acharya and Menon [1] discuss many examples in their review of segmentation methods. The advantage of using only simple primitives is they incorporate very few assumptions about the objects represented and are thus very flexible. The most common problem with these techniques is instability. For example, edge-based approaches depend strongly on the characteristics of the edge detector used.

The most popular and generally effective group of shape models used for feature extraction are *deformable models* [6, 15]. Examples include energy minimizing *snakes* [5, 7] and 3D *deformable surfaces* [16, 8]. A limitation of most of these techniques is that they have to impose smoothness constraints on the models, which limits the types of objects that can be represented with any given set of parameters. This may be problematic especially if little or no *a priori* information about the image is known, such as in many exploratory applications.

Another group of popular shape models are based on a medial axis or skeleton representation (*e.g.*, [4, 2, 3]). One of the most comprehensive medial axis approaches to date is by Pizer *et al.*, who propose a shape model that can be used for various applications in medical image processing [10] and computer graphics [11]. Pizer’s approach represents 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. One of the significant shortcomings of the Pizer model is that the shape metric only works with objects having the same number of figures and medial primitives, which strongly limits its flexibility.

The algorithm presented in this paper is based on the UoC representation pioneered by Ranjan and Fournier [14]. UoC models have a number of desirable properties [12] that make them preferable over many of the models mentioned above for feature extraction and subsequent shape-related processing. For the application in this paper, the most relevant properties are stability, flexibility and a shape metric that can be used to quantify the difference between any two UoC's regardless of their differences in topology or number of components.

3 Methodology

This section describes the details of our algorithm for extracting shape models from images. The simple image shown in Figure 4 is used as an example to explain the processes involved. This image (150×150, 8 bits per pixel) contains two objects, one having a uniform intensity (100), the other having a linear intensity gradient ($175 - 225$).

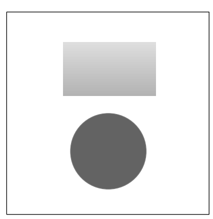

Figure 4: Simple test image

As illustrated in Figure 5, the main steps of our algorithm are:

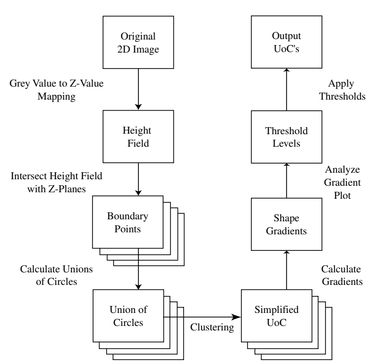

Figure 5: Overview of shape model extraction algorithm

1. Generate boundary points.
   i. Generate a height field from the image's intensity data by mapping the value at each pixel to a z-value.
   ii. 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 UoC for each set of boundary points.
3. Simplify each UoC by clustering.
4. Calculate the shape gradient between successive levels using the UoC method of shape matching.
5. Determine the thresholds of interest from the maxima and minima in the shape gradient data.
6. Use the UoC's from Step 3 and the thresholds from Step 5 to compute UoC's of the objects in the image.

3.1 Boundary Point Generation

Creating a height field from a greyscale image is relatively straight-forward. All that is required is a mapping from the intensity value at each pixel to a z-value in the height field. The simplest case would be a linear mapping. The points in the height field are then used as the boundary points of a volume that is bounded below by the image plane. Figure 6a shows the volume created from the example in Figure 4 using a linear mapping.

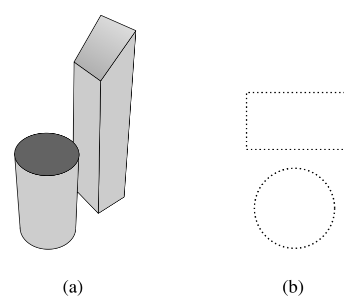

Figure 6: (a) Volume created from the test image (b) Boundary points created from intersecting the volume with a z-plane

The resulting volume is then intersected with n different z-planes. The values used for z are $\{I_i, i = 1, \ldots, n : I_{min} \leq I_i \leq I_{max}, I_{i+1} = I_i + I_{incr}\}$, where I_{min} and I_{max} are the minimum and maximum intensity values of the image, and I_{incr} is the intensity increment from one level to the next. The intersection of each plane with the volume results in a set of boundary points at that level. Figure 6b shows the boundary points created by intersecting the test volume with the plane $z = 50$.

3.2 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 [13, 14]):

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 7a shows the UoC generated from the boundary points in Figure 6b. This model has 358 circles.

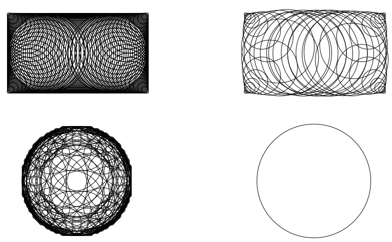

(a) (b)

Figure 7: (a) UoC (358 circles) generated by intersecting the volume in Figure 6a with the plane $z = 50$ (b) Simplified version of the model, with only 46 circles

3.3 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.

Figure 7b shows a simplified version of the UoC in Figure 7a. The original UoC has 358 circles, with 150 circles representing the disk. 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. Using dense points ensures that no features are missed. After simplification, the resulting UoC, with only 46 circles, is a more efficient representation and is more useful for comparison with other UoC's.

The simplification algorithm, called *clustering*, works by taking clusters of circles within the UoC, and replacing each cluster with a single encompassing circle. 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 [13, 14].)

3.4 Shape Gradient Computation

The next 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 8 shows three UoC's computed from the image in Figure 4 at three different levels.

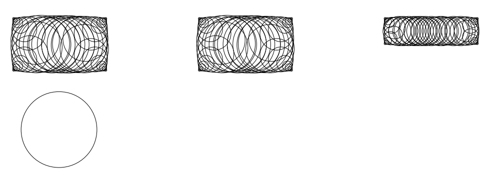

(a) (b) (c)

Figure 8: UoC of Figure 4 at three intensity levels (a) 50 (b) 125 (c) 200

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 neighbours. Each circle has a maximum of three neighbours because each triangle has a maximum of three neighbours. Between the circle and each neighbour, we take the gradient $\frac{dR}{dD}$, where dR is the signed difference between the radius of the circle and the radius of the neighbour, and dD is the Euclidean distance between the centres of the circles. The gradients in the directions of the three largest neighbours form the feature of the circle. The feature distance between two circles can be best explained using a physical analogy. If the two features have a common centre and are free to rotate around it, and between the ends of each pair of gradient vectors (in the directions of the neighbours) there is a spring that has a pulling force proportional to its length, then the system will be at rest when the potential energy is at a minimum. The sum of the residual distances between the vector ends in this minimum energy state is taken to be the feature distance between the two circles. Thus, our cost function defines 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 (*e.g.*, 1.0) to each weight and use the same three values for all images.

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 maximal 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 maximal 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) 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 images (such as in the medical domain), the matched gradient plot alone is sufficient for threshold determination, the unmatched gradient plot is sometimes useful for providing complementary information.

3.5 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 spans a relatively large intensity range; a sharp spike means the object spans a narrow intensity range.

Figure 9 shows the shape gradient plot computed from the image in Figure 4. The matched gradient clearly shows a significant object, in this case the rectangle, in the range $175 \leq I \leq 225$. The unmatched gradient in this range confirms the result. The matched gradient increases here because as the rectangle gets thinner, its shape change becomes more significant relative to its size. The unmatched gradient decreases because the unmatched circles are getting smaller. The unmatched gradient also shows a peak at $I = 100$. This means there is an object of uniform intensity, in this case the disk, at that level. The height of the peak indicates that the object is of significant size.

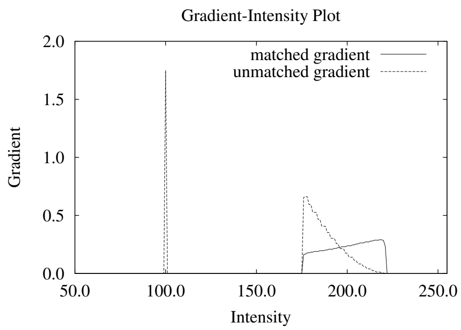

Figure 9: Shape gradient plot computed from the image in Figure 4

In this example, the maxima and minima are very obvious and distinct from each other. However, this is often not the case with real images. This brings into question how we should decide which maxima and minima are to be considered significant. These decisions currently require user input and are based on the level of detail at which the user wishes to analyze the image. Applying a Gaussian filter to the plot to smooth out small maxima and minima is a helpful step.

3.6 Shape Model Generation

Having determined the upper and lower thresholds of a significant object, we can compute a shape model for that object by using the UoC's generated during the gradient computation process. As demonstrated in the Results section, for some applications it is best to simply use the UoC at the lower threshold as the object's shape model. In other cases, we take the lower threshold UoC, and remove from it any circles or parts of circles that would lie within the upper threshold UoC if the two UoC's were superimposed. If the resulting UoC has partial circles in its boundary, it is retriangulated (Section 3.2) to form a new UoC. The result is a UoC of the object in the appropriate intensity range.

4 Results

This section presents the results of applying our technique to three representative test cases, all of which are actual medical images. The quality of these results shows that the method is potentially viable in practical applications. For all of the test cases, we use a linear mapping for the height field and an increment of 1.0 between threshold levels. For simplification, 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 CT and MRI images. The three weights w_p, w_s and w_f in the cost function of the distance metric are set to 1.0 for the shape gradient calculation.

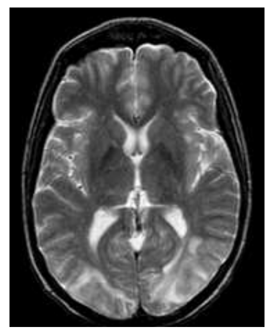

Figure 10: Brain MRI

The first test case, shown in Figure 10, is an MRI image of a human brain. This image has 128 grey levels. Figure 11 shows the shape gradient plot computed from this image. Using the analysis method outlined in Section 3.5, we can clearly

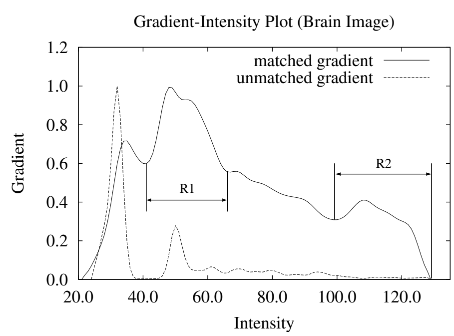

Figure 11: Shape gradient plot computed from the brain MRI image

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 shape analysis. The other two ranges, labelled R1 and R2 in Figure 11, are for the whole brain and the grey matter inside the brain.

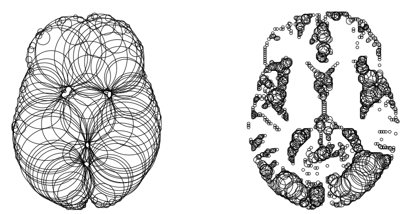

Figure 12: UoC's computed from the brain MRI at two levels (a) 43 (whole brain) (b) 100 (grey matter)

Figure 12 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 shape stability at these intensities, which are the lower thresholds for the whole brain and the grey matter inside. In this case, the UoC's for the upper thresholds are not used. The reason we do not use the upper threshold for the whole brain is that the UoC at $I = 43$ is a good representation of the shape of the brain as a solid mass, and this is likely to be of more use in the analysis of the overall shape of the brain than the same model with parts of its interior taken out. We do not need to use an upper threshold for the grey matter, because R2 is at the top of the image's intensity range. Therefore, we use the UoC's at the lower thresholds (Figure 12) as the final shape models for the two objects.

The second test case, shown in Figure 13, 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.

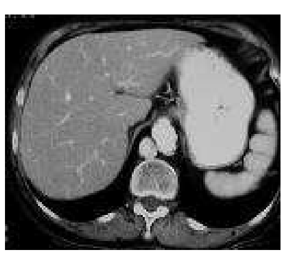

Figure 13: Lower abdomen CT

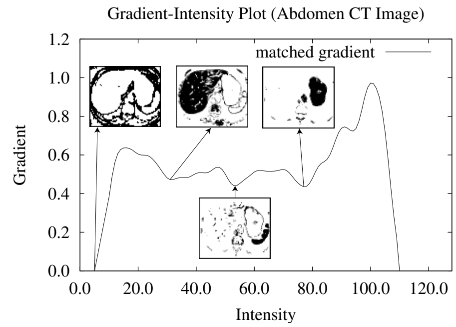

Figure 14: Shape gradient plot computed from the abdomen CT image

Figure 14 shows the matched gradient plot from this 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 14 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. However, there is not a distinguishable maximum for the spine, which points out a weakness of the method. The intensity range of the spine overlaps with those of the liver and small intestine, and because the spine is significantly smaller than the other two structures, it is essentially lost. This kind of problem is common to all algorithms using global thresholds.

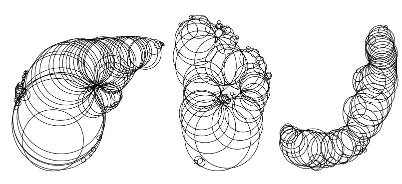

Figure 15: (a) UoC of liver (b) UoC of kidney (c) UoC of small intestine

Figure 15 shows the final shape models of the liver, kidney and small intestine computed using their respective upper and lower thresholds, with some manual removal of circles not connected with the objects.

The third example, shown in Figure 17a, 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 methods such as histogram-based analysis very difficult.

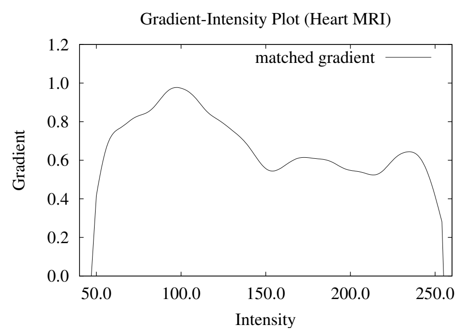

Figure 16: Shape gradient plot computed from the first frame of the heartbeat MRI sequence

Figure 16 shows the shape gradient plot for the first frame of the sequence. 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 threshold values for the first object range from 48 to 51, the values for the second object from 120 to 172 and the values for the third object from 178 to 226.

Figure 17b 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 suitable thresholds.

5 Summary and Future Work

We have presented a novel method for exploring shape information in a greyscale image. Our algorithm uses the shape gradient to effectively extract thresholds and shape models of significant objects. The Union of Circles model is used to compute the gradient and represent the extracted shapes. We presented the test results of applying our algorithm to artificial images and real medical images. These results indicate that the our algorithm is potentially viable in practical applications. The method presented is very new and there are many directions for future research. At the top of the list are automatic methods for accurately locating points of interest on the gradient curve and more formal and stringent validation of the results of our algorithm.

6 Acknowledgments

This paper is dedicated to the memory of Alain Fournier. I would like to thank Wolfgang Heidrich, Rob Walker, Rob Scharein, Lisa Streit and the VMV reviewers for their insightful comments. Funding for this research is provided by NSERC and the University of British Columbia.

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] D. Attali and A. Montanvert. Semicontinuous skeletons of 2D and 3D shapes. In *Proceedings of the International Workshop on Visual Form*, pages 32–41, Capri, 1994. World Scientific.

[3] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding*, 67(3):161–273, 1997.

[4] 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.

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

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

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

[8] T. McInerney and D. Terzopoulos. T-Snakes: Topology adaptive snakes. *Medical Image Analysis*, 4(2):73–91, June 2000.

[9] 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.

[10] 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.

[11] S.M. Pizer, A.L. Thall, and D.T. Chen. M-reps: A new object representation for graphics. Technical Report TR99-030, University of North Carolina at Chapel Hill, September 1999.

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

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

[14] 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.

[15] A. Singh, D. Goldgof, and D. Terzopoulos, editors. *Deformable Models in Medical Image Analysis*. IEEE Press, Los Alamitos, CA, 1998.

[16] 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.

Frame	Heart Walls	Chambers	Blood
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: Lower thresholds extracted from heart MRI frames

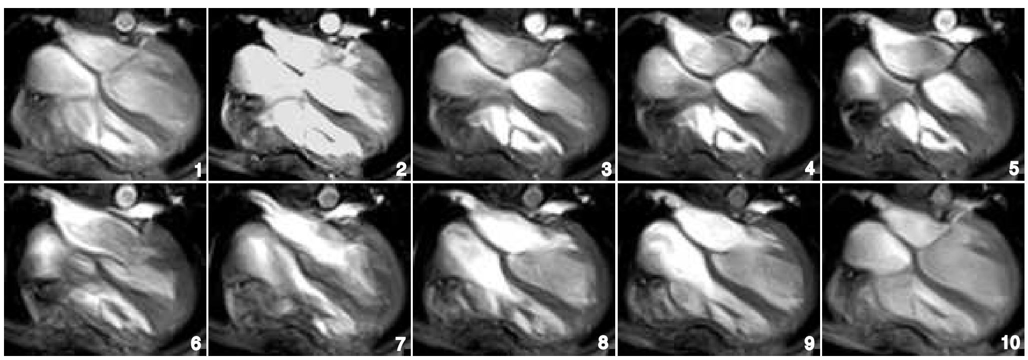

(a)

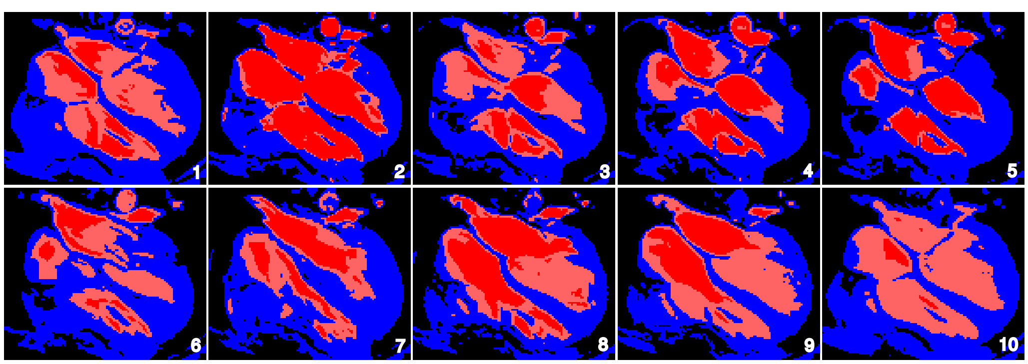

(b)

Figure 17: (a) Time-gated MRI sequence of a heart during a single cardiac cycle (b) Colourized version of the same sequence using shape gradients for threshold extraction