

Segmentation of the Radiation Treatment Field in Dual Portal Images

by

Bryce Cutt

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Bachelor of Science Honours

in

The I. K. Barber School of Arts & Sciences

(Computer Science)

The University Of British Columbia

April, 2007

© Bryce Cutt 2007

Abstract

Conformal radiotherapy involves many treatments and multiple beams of radiation applied to the patient. For conformal radiotherapy the successful detection of the radiation field resulting from a treatment is critical for quality assurance purposes. Errors over multiple treatments can compound resulting in the death of healthy cells while allowing cancerous cells to survive. Often the verification is performed manually by a technician which, across multiple treatments, is very time consuming.

However, automated methods have been developed using gradient thresholds to segment the dual portal image for field detection. I present an improvement of the accuracy of the generic gradient threshold technique achieved by incorporating the gradient directional information. I have implemented other segmentations techniques, such as the Fast Marching Method, that propagate a curve based on the properties of the image with the final curve location representing the detected field edge. Though the directional gradient threshold method is often more accurate than previous methods, it is less easily automated; whereas the curve propagation methods are easy to automate, are less computationally intense, and produce more accurate segmentations.

Acknowledgements

I would like to thank my supervisors Dr. Patricia Lasserre and Dr. Yves Lucet for providing me with the opportunity to do an Honours under them. Their feedback and insight has been invaluable in bringing my thesis to its current form. I would also like to thank Dr. Lasserre and Dr. Lucet for mentoring me throughout my degree and providing many opportunities for me to build the confidence needed to pursue an Honours thesis.

Dr. Rasika Rajapakshe of the British Columbia Cancer Agency provided the images and his medical expertise and for that I am very grateful.

Many people have provided feedback and support regarding my thesis and I would like to express my appreciation to them and acknowledge their contribution to my success.

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Figures	v
List of Tables	vi
1 Introduction	1
1.1 Radiotherapy Treatment	1
1.2 Double Portal Images	2
1.3 Image Segmentation	5
2 Edge Detection Techniques	6
2.1 Gradient Calculation	6
2.2 Directional Gradient	7
2.3 Curve Propagation	13
2.3.1 Level Set Method	13
2.3.2 Fast Marching Method	14
2.3.3 Fast Sweeping Method	14
2.3.4 Segmentation Using Curve Propagation	15
3 Results and Conclusions	18
3.1 Choice of Curve Propagation Method	18
3.2 Comparative Results	19
3.3 Future Work	21
Bibliography	22

List of Figures

1.1	Three-dimensional conformal radiotherapy treatment machine.	2
1.2	A dual portal image showing the targeted area in dark, and the patient internal organs and bones.	3
1.3	Complex images	4
2.1	Comparison of a double portal image and its morphological gradient	6
2.2	Directional gradient computed using the center of the picture as the center point	7
2.3	Comparison of gradient method cross sections. Notice that the directional gradient has far less noise than the standard gradient.	8
2.4	Directional gradient of a complex image computed using the center of the picture as the center point	9
2.5	Directional gradient of complex image using user selected center points	10
2.6	Resulting segmentations of a complex image by the Wang-Fallone method using the directional gradient	11
2.7	The vectors involved in a directional gradient. The center of the image is indicated by a black dot. The gradient direction is indicated for some edges by a dashed line and the vector from the center of the image is indicated by a solid line.	12
2.8	Curve propagation segmentation method using initial points	16
2.9	Curve propagation segmentation method using initial outline	17
3.1	Comparison of segmentation results using each method	20

List of Tables

3.1	Comparative results	19
3.2	Comparative computation time	21

Chapter 1

Introduction

An estimated 72,700 people will die of cancer in Canada in 2007. On the basis of current incidence rates in Canada, 38% of women and 44% of men will develop cancer during their lifetime, while 24% of women and 29% of men will die from cancer [1]. It is estimated that approximately 30% of cancer treatments require a combination of surgery and radiotherapy. The goal of these treatments is to increase patient survival rates and minimize damage to healthy organs during treatment.

1.1 Radiotherapy Treatment

Three-dimensional conformal radiotherapy is a method for the treatment of cancer. A treatment plan is created that will apply the proper amount of radiation to the cancerous cells while minimizing damage to healthy cells. The plan indicates the shape, intensity, and location of multiple beams of radiation that will traverse the patient's body during the treatment. If some beams are incorrectly shaped, or the location they are applied to is faulty by even the smallest amounts, the radiation will be applied to the wrong areas, which can potentially kill the healthy cells while allowing cancerous cells to live. It is important that these errors are detected before they can compound and significantly deteriorate the health of the patient.

Currently a technician has to manually compare the radiation applied during treatment to that called for by the plan. If this process can be automated, the technician can spend more time treating the patient and potentially allow more patients to benefit from such treatments.

A linear accelerator, such as that shown in Figure 1.1(a), is used to treat a patient. The patient lies on the bed and the treatment machine rotates to the correct location. Radiation is then passed through the patient's body and measured by the detector shown at the bottom of Figure 1.1(a). The machine produces a double portal image, using the detected radiation, as described in Section 1.2.

The shape of radiation produced is controlled by blocking part of the radiation beam. This is accomplished by using a block where the cut out portion is the desired shape of the radiation or by using a multi-leaf collimator as shown in Figure 1.1(b). Each leaf of the multi-leaf collimator can move independently in or out so that the correct shape of radiation is produced by the machine.



(a) The radiation detector and patients bed. The detector can rotate 360 degrees around the patient so that radiation can be passed through at the appropriate angle.



(b) Multi-leaf collimators are often used to shape the treatment radiation.

Figure 1.1: Three-dimensional conformal radiotherapy treatment machine.

1.2 Double Portal Images

A portal image is a monochromatic image in which the intensity of each pixel is inversely proportional to the intensity of radiation detected in that portion of the sensor. A double portal image involves two exposures of radiation recorded in the same image. The first exposure is a very light radiation that shows the anatomy of the patient and the orientation of their body. The second exposure is the radiation used to treat the cancer. The final image resembles a picture of the anatomy of the patient, with darker patches where the patient's body did not block much of the radiation and at the location of the treated area (usually around the center of the image, see Figure 1.2).

The image set used for testing was generated by two different machines resulting in images of different sizes and quality. Some of the images are 16-bit monochromatic images that are 1024 by 1024 pixels in size (such as Figure 1.2) and are generated by the Elekta SL 20 linear accelerator which uses an iView GT portal imager. The remaining images are 8-bit monochromatic and 768 by 576 pixels in size (such as Figures 1.3(a) and 1.3(b)) which are generated by a SL-75 linacs video system. Images of varying complexity were available for testing the edge detection techniques discussed in this paper so a sample that represented a range of complexities and areas of the body was selected with a bias towards complex images that would prove challenging to segmentation techniques.


Of the 65 images in the image set, 32 are complex. A complex image is one that standard




Figure 1.2: A dual portal image showing the targeted area in dark, and the patient internal organs and bones.

segmentation techniques tend to fail miserably on. Images of this type usually have weak, gradual, or broken edges that cannot be easily detected as a proper edge. These complex images could also have strong internal edges that cause over segmentation (most standard segmentation techniques do not handle large internal edges very well). Internal edges, and often edges that divide the inside of the true radiation field, are quite common in head and neck images as the actual radiation field is often located partially inside the patient and partially outside the patient.

Figure 1.2 is a good example of a relatively simple image where the treatment radiation is located entirely within the patient and there are only small variations and edges within the treatment field. Upon close inspection it is apparent that the edges of the treatment field are not sharp edges but are instead gradual across multiple pixels. While this image is classified as simple, the gradual edges do make it too complex for the standard segmentation techniques. Figures 1.3(a) and 1.3(b) are complex images. Both have edges inside the treatment field that are as strong or stronger than the edges of the treatment field. Due to the patient's body blocking portions of the treatment radiation (and in the case of Figure 1.3(a) the treatment field falling partially outside the patient) there are large variations in the intensity of the pixels that form the treatment field in the image. Other complicating factors are the very gradual edges at the top and middle bottom of Figure 1.3(a), and the top right of Figure 1.3(b), at the edge of the treatment field.



(a) Complex mandible image due to internal edges and separation of field



(b) Complex lung image due to gradual edges and internal edges

Figure 1.3: Complex images

1.3 Image Segmentation

Image segmentation is the process of separating an image into multiple distinct, non-overlapping parts. In many cases, including double portal images, the image is to be separated into two parts. The area of interest in a double portal image is the radiation treatment field, as described in Section 1.2, and the goal of segmenting the image is to determine exactly which pixels in the image are within this field and which are not.


Many edge-detection techniques have been tested by previous investigators to find the edge of the radiation treatment field. Gilhuijs and van Herk used a gradient for histogram thresholding of portal images [2, 5]. A wavelet based, multi-scale edge detection technique was presented by Petrascu et al [15]. Graph-based algorithms [11], watershed methods [7, 12], and deformable models have been used extensively in other areas of medical imaging. The Level Set Method, Fast Marching Method, and other deformable model methods have been used to segment many types of medical images but have not yet been used for double portal images to the best of our knowledge [3, 4, 8, 9, 10, 14, 16, 17, 18, 19, 21].

Beyond the implementation of standard segmentation techniques, the first advanced edge detection technique used for double portal images was a gradient threshold algorithm proposed by Wang and Fallone [22, 23]. The method involves first taking the gradient of the image, and then testing a series of thresholds of that gradient until an optimal one is found. The basic idea of an optimal threshold hinges on the fact that all final segmentations should contain a minimum number of large shapes in them (treatment area, image border area, human body area). The number of large shapes can be counted at each sequential threshold tested and the lowest threshold with the minimum number of shapes is selected as the optimal threshold. Once the optimal threshold is found, the central shape is selected (as being the shape that contains the central pixel of the image) and its outline is used as the outline of the final segmentation. Based on the initial sample of images, there are always three large shapes of at least a fixed size in each properly segmented portal image from the Elekta SL 20 linear accelerator, and often only two in images from the SL-75 linacs video system.

Chapter 2

Edge Detection Techniques

2.1 Gradient Calculation



(a) Cross section of the double portal image


(b) Cross section of the morphological gradient image. (The intensity has been scaled up and the inverse image is shown for better visibility).

Figure 2.1: Comparison of a double portal image and its morphological gradient


A gradient image represents the change in an image. The change is measured for each pixel by calculating the difference between a pixel and its neighboring pixels. The intensity of the gradient is directly proportional to the change in intensity of the original image. This calculation can be seen as the absolute first derivative of a three dimensional function with respect to intensity where the location of the pixel in x and y and the intensity are the three dimensions. Figure 2.1(a) is a cross section of the dual portal image. Note that the transition between dark and light is represented by a slope (and change in height) in the cross section. Also note that the edges, especially the edges of the treatment field, are gradual and spread among multiple pixels. Figure 2.1(b) is the gradient image. In this image the darker the pixel (and higher the cross section) the larger the gradient for that pixel. It can be seen that



(a) The treatment field edge is a transition from dark to light in the direction away from the field



(b) Regular morphological gradient. The edge inside the treatment field is part of gradient. (The inverse image is shown for better visibility).



(c) Directional morphological gradient. Edges inside the treatment field are discarded along with some of the external edges that would cause problems with the automatic thresholding. (The inverse image is shown for better visibility).


Figure 2.2: Directional gradient computed using the center of the picture as the center point

the height of the cross section in Figure 2.1(b) is directly proportional to the steepness of the slope in Figure 2.1(a). The cross sections display quite clearly that there is a large amount of minor noise in the double portal images even though it is difficult to see without the cross section. It is also clear that the treatment field can contain a large range of intensities; it is not consistent across the entire field even in areas that look consistent.

There are many ways to perform a gradient calculation. Two classic ones are the Sobel gradient and the Prewitt gradient [6]. Essentially, any calculation that results in the absolute change across a pixel can be used, such as subtracting the intensity of the previous pixel from the intensity of the next pixel in both the x and y directions and using the magnitude of this two dimensional vector as the pixel intensity in the gradient. The morphological gradient is calculated by subtracting the minimum of the image (the erosion) from the maximum of the image (the dilation) [6]. This method produces an accurate gradient very quickly, hence the reason it is recommended by Wang and Fallone [22] for use in their gradient threshold algorithm.

2.2 Directional Gradient

The classic computation of a gradient for image segmentation only takes into account the gradient magnitude at each pixel; the direction of the gradient is either not calculated or discarded. An important tool for image segmentation is to integrate any information that is common among all images being segmented into the algorithm to enhance the accuracy of the segmentation. For dual portal images from the two treatment machines the area that




(a) Cross section of a standard morphological gradient image. (The inverse image is shown for better visibility).

(b) Cross section of the directional morphological gradient image. (The inverse image is shown for better visibility).

Figure 2.3: Comparison of gradient method cross sections. Notice that the directional gradient has far less noise than the standard gradient.




(a) Complex image due to hip replacement




(b) Directional gradient from center of image where hip replacement edge is not thrown out as needed. (The inverse image is shown for better visibility).

Figure 2.4: Directional gradient of a complex image computed using the center of the picture as the center point




(a) Center points for the directional gradient are manually chosen



(b) Hip replacement edge is throw out due to better center points. (The inverse image is shown for better visibility).

Figure 2.5: Directional gradient of complex image using user selected center points



(a) Final segmentation using the center of the image as the center point (b) Final segmentation with user selected center points

Figure 2.6: Resulting segmentations of a complex image by the Wang-Fallone method using the directional gradient

must be segmented from the rest of the image is always darker than the surrounding parts of the image which means that the edge to be detected is always a transition from dark to light in the direction away from the shape to be segmented (see Figure 2.2(a)).

Integrating the direction of the gradient resulted in a new gradient computation that allows all edges that do not transition from dark to light in the direction away from the treatment field to be discarded. The first version of the directional gradient method (illustrated in Figure 2.7) would take the vector from the center of the image (indicated by the solid arrow) to the point under inspection and compare it to the gradient vector (indicated by the dotted arrows). If they were in the same direction (within 90 degrees) the gradient would be kept, if not then the gradient would be discarded. In Figure 2.7 it can be seen that all parts of the proper field edge are in the same direction as the vector from the center of the image while all parts of the internal edge are in the wrong direction. The result of this method is illustrated by comparing the regular morphological gradient (Figure 2.2(b)) and the directional morphological gradient (Figure 2.2(c)) of a dual portal image (Figure 2.2(a)).

Figure 2.3(a) and Figure 2.3(b) show cross sections of the gradients in Figure 2.2(b) and Figure 2.2(c) respectively. As can be seen by comparing the two images, the strong edge inside the treatment field, two of the stronger edges outside the treatment field, and much of the noise are thrown away properly because they are a transition in the wrong direction. On the contrary, the edges we want to keep are kept in their entirety with no weakening of their gradient.

In many images, using the vector from the center point works well. However, due to




Figure 2.7: The vectors involved in a directional gradient. The center of the image is indicated by a black dot. The gradient direction is indicated for some edges by a dashed line and the vector from the center of the image is indicated by a solid line.

the complexity and variety of shapes, the vector from the center point may not always be pointing in the same direction as the edge of the treatment field in some parts of the image. If there is a large concavity in the treatment field edge it is possible for part of the edge to be a transition from dark to light in the opposite direction of the vector from the center point. It is also possible for an internal edge of the image to be a transition from dark to light in the direction away from the center point, which will remain strong in the resulting directional gradient. This is the case with Figure 2.4(a), where the right edge of the hip replacement is a transition from dark to light in the direction away from the center of the image. In its directional gradient (shown in Figure 2.4(b)) the left edge of the hip replacement is discarded but the right edge remains, and separates the treatment field into two parts thus defeating the automatic thresholding method and resulting in the final segmentation shown in Figure 2.6(a). Figure 2.4(a) demonstrates that for the directional gradient to be successful

on the range of images under consideration the center of the image cannot always be used to determine the direction of gradients to keep. The solution is to allow the user to choose one or more center points. This requires user input and can become both quite complex and time consuming. In the case of the image in Figure 2.4(a) a good working set of points are shown as crosses in Figure 2.5(a) resulting in the directional gradient shown in Figure 2.5(b) and the final segmentation shown in Figure 2.6(b). Finding this appropriate set of points took more time than to draw the entire proper outline by hand.

2.3 Curve Propagation

A modern image segmentation technique involves modeling the segmentation as the propagation of a curve that, upon completion of the algorithm, separates the area to be segmented from the rest of the image. This technique, often referred to as active contours, snakes, or parametric contours was first presented by Kass et al. [9] and many variations have been developed by other authors [4, 10]. Some of the more recent curve propagation methods are the Level Set Method [14, 18, 19, 21], the Fast Marching Method [17], and the Fast Sweeping Method [20].

2.3.1 Level Set Method

The level set method, when used for image segmentation, works on the principle that a curve can propagate within an image by iteratively applying forces to the curve. The forces applied are an expansion force that pushes the curve outward, an attraction force that pulls the curve towards areas of interest, and a curvature force that causes the curve to have a smooth curvature with no sharp points. The level set method (like its variations) can be used for many applications such as physics modeling and computer animation. For image segmentation it can generate a smooth segmentation outline that successfully attracts itself to strong gradient edges in the image. In image segmentation, the attraction force is usually based on the inverse of the image gradient so that a strong gradient has an almost zero attraction force, causing no movement of the curve.

For each iteration of the level set method the new position of the curve is calculated. To minimize the complexity, calculations are only done on a subset of pixels around the current curve. The subset consists of all pixels within a set distance from the current curve which is referred to as the narrow band. The curve can move in any direction and any distance within the narrow band at each iteration depending on the features of the image and the forces being applied (it can break into multiple shapes, expand, and contract). Over the course of multiple iterations the method can adjust itself to quite accurately match the desired shape if the correct parameters are used.

2.3.2 Fast Marching Method

The Fast Marching Method is a special case of the Level Set Method that can only move in one direction (expanding or contracting). Like the level set method the Fast Marching Method is subject to three forces but the attraction force and curvature force cannot force the Fast Marching Method to move backwards. For the Fast Marching Method to be used for image segmentation it must start from an initial point, or set of points, and it must have a direction. The propagation of the method is modeled by taking the current location of the curve and calculating the distance from the curve to all adjacent points in the direction that the curve is moving. Each iteration, the curve moves to expand over the next closest pixel of the image and the distance to all adjacent pixels is calculated.

This form of propagation to adjacent pixels is much like the way the narrow band is used in the Level Set Method except that it is only ever on one side of the curve and it is always one pixel in width. The direction is modeled by applying a constant speed. The attraction force is calculated from the gradient at the pixel in question. The combination of constant speed and gradient makes the distance between two directly adjacent pixels equivalent to the constant speed when the gradient is zero. The distance is much higher than the constant speed when the gradient is high. Remember that the Fast Marching Method expands over closer pixels first so if the gradient is high (making the distance large) then the method will take a long time to expand over that pixel. This makes the fast marching method move very quickly across parts of an image that are of consistent pixel intensity (zero gradient) but more slowly, or not at all, at parts of an image that have a large gradient, such as the edges the method is meant to find. For more information on implementing the Fast Marching Method, see the work by Lin [13] and Sethian [18].

2.3.3 Fast Sweeping Method


Much like the Fast Marching Method, the Fast Sweeping Method [20] is a restricted case of the Level Set Method that can only move in one direction. While the Fast Marching Method starts from a point, or set of points, and then expands one pixel at a time until reaching the final segmentation, the Fast Sweeping Method passes over the entire image multiple times, from corner to corner, updating the distance values for all pixels on each pass. After a set number of passes all pixels within a given distance from the initial shape are treated as inside the shape and all others outside, thus giving the outline of the shape as the border between these two sets of pixels. The Fast Marching Method's computational time is largely dominated by sorting the set of adjacent pixels so that the next closest one may be chosen and added to the shape. The Fast Sweeping Method was developed to get around the sorting and instead just pass over all pixels iteratively as described above. This reduces the complexity of the algorithm from the $O(N \log N)$ of the Fast Marching Method to $O(N)$ for the Fast Sweeping Method where N is the number of pixels in the image. Based on solely the complexities, the Fast Sweeping Method is much faster than the Fast Marching Method for large images.

2.3.4 Segmentation Using Curve Propagation


Ideally a segmentation using a curve propagation method should start from an initial point and expand to the border of the shape to be segmented. For segmenting the image in Figure 2.8(a) the center of the image could be used as the starting point as this would require no input from the user to select initial points. However, as can be seen in Figure 2.8(b) the expanding curve will break out of the edges of the treatment field before it will expand over the strong internal edge in this image. A solution may be to try using multiple starting points like in Figure 2.8(c) or even more points such as in Figure 2.8(d) but this adds lots of user input to select each point and a good set of points is difficult to choose as any relatively strong internal edge, such as that separating the top of the treatment field or the left side of the treatment field in Figure 2.8(a), would require many points to be manually and carefully placed along each internal edge.

Some additional information available for each image to be segmented is the plan image that represents the shape and location of radiation that was intended to hit the patient (See Figure 2.9(a)). If an outline is generated from this plan image, it can be used so that all points within the outline are treated as initial points for the curve propagation method. Since the intent of this method is to detect what radiation actually hit the patient and considering the Fast Marching Method and Fast Sweeping Method can only expand or contract, the generated outline cannot simply be used as is. Instead, the method must guarantee that the outline is strictly within the treatment field by a margin greater than the possible error in the treatment field. This task is done by taking the outline generated from the plan image and eroding it a few times so that the initial outline used is smaller than the generated outline. The user then clicks on the treatment image to place the initial outline inside the treatment field such as in Figure 2.9(b) and the curve propagation method runs until the stopping condition is reached resulting in a final segmentation such as that shown in Figure 2.9(c).


There are many possible stop conditions for a curve propagation method but it was discovered through experimentation that an effective stop condition for this application is to allow the curve propagation method to expand a set distance from the initial points. Since the distance between two points is largely based on the gradient along the path between those two points and the gradient magnitude can be much larger for higher bit rate images it makes sense that there is a different good distance for the stop condition for 8-bit images than for 16-bit images.




(a) The original image



(b) Segmented using the image center point as the starting point




(c) Segmented using multiple points as starting points




(d) Segmented using many points as starting points


Figure 2.8: Curve propagation segmentation method using initial points



(a) The plan image used to generate the initial outline



(b) The initial outline generated from the plan image and placed inside the treatment field to be segmented



(c) The final segmentation

Figure 2.9: Curve propagation segmentation method using initial outline

Chapter 3

Results and Conclusions

3.1 Choice of Curve Propagation Method

Despite calculation only being performed on the narrow band, the Level Set Method is very computationally intense. On top of being slower than the Fast Marching Method and Fast Sweeping Method, it is harder to find the correct parameters for the method to be successful in a particular application. The main parameters to be determined are the number of iterations and the strength constants for the three forces that move the curve. For segmenting double portal images a successful set of values for the force constants has not been found so the Level Set Method has not been used successfully so far.

Since the resulting distance values for the image are very close using both the Fast Marching Method and the Fast Sweeping Method, it would seem that the best choice for the curve propagation segmentation would be the Fast Sweeping Method due to its speed. What has been discovered, however, is that the Fast Marching Method only has to pass over the pixels in the image that are to be segmented out; in the double portal images that number of pixels is always less than half the number of pixels in the image. So, effectively, the Fast Marching Method is always faster for the double portal images. The recommended minimum number of passes over the image for the Fast Sweeping Method (and the number that has produced relatively good results) is four times; once from each corner. With big O notation, constant multiples in the runtime are discarded but in real world examples when the size of N is not infinitely large those constant multiples can come into effect.

The Fast Marching Method's complexity is $O(N \log N)$ because in the worst case, for every N pixels, the algorithm must sort a set of pixels that could include all N pixels to find the next one to expand over. In practice, the first N pixels are the number of pixels the method expands over and since it does not have to expand over all pixels in the image for double portal images that number is always less than N. The $\log N$ portion assumes the worst case, where the method must sort all pixels for every pixel, but this is not the case. In practice the method must sort the narrow band (all pixels adjacent to the current curve) for every pixel to be expanded over and while the narrow band constantly increases in size it is never anywhere near as large as N.

In contrast the Fast Sweeping Method is $O(N)$ and must sweep over all N pixels multiple times for it to be effective. In practice the Fast Sweeping Method makes four passes over the image. Images in the image set have a maximum size of 1024 by 1024 pixels.

Because the Fast Marching Method only has to visit the pixels that are to be segmented out the difference in complexity of the two algorithms does not decide which is faster for this application. For our application, the Fast Marching Method performed significantly faster

Table 3.1: Comparative results

Segmentation Method	Image Set	Results Acceptable		
		Easy	Medium	Complex
Wang-Fallone	16-bit images	18 of 18	2 of 4	0 of 3
Fast Marching Method	all images	19 of 19	14 of 14	31 of 32

than the fast Sweeping Method. Obviously, for significantly larger images the two algorithms would have to be re-evaluated.

3.2 Comparative Results

The Wang-Fallone method has been successful on some of the simpler images in the image set but it fails miserably on any of the more complex images that contain strong internal edges or gradual edges. This can be seen in Figure 3.1(b), which shows the Wang-Fallone method using a standard morphological gradient applied to the image in Figure 3.1(a). With great care in choosing center points the Wang-Fallone method using the directional morphological gradient is able to do a very good job segmenting out the proper treatment field. Unfortunately this is not the case for complex images. The method requires a large amount of time for the user to find a good set of center points. A good set of center points for most of the complex images, resulting in segmentations such as Figure 3.1(c), have not been found. In any case, the resulting segmentation is incomplete and of unacceptable quality.

Part of the ease in automating the Wang-Fallone method is that there is no variation in parameters such as the number of large shapes in an optimal threshold. After expanding the image set it is clear that not all of the 8-bit images can be split by an optimal threshold into the same three shapes as the 16-bit images can. It was also discovered that a small percentage of images in the expanded image set have a treatment field that does not fall on top of the center of the image and that, in fact, the treatment field may be composed of multiple disconnected shapes. These two facts prevent the method from being easily automated as the user would need to provide specific information to the method for each image. Of the images in the image set that the Wang-Fallone method can be automated on, eighty percent are properly segmented; but this is less than thirty-one percent of the total number of images in the set (see Table 3.1).

In contrast with the Wang-Fallone method, the curve propagation method properly segments more than ninety-eight percent of the images in the image set (see Table 3.1). The curve propagation method also produces much better segmentations than the Wang-Fallone method as can be seen by comparing the edges in Figure 3.1(c) to the same edges segmented by the curve propagation method in Figure 3.1(d). The curve propagation method produces smoother, more accurate edges.

On top of the accuracy difference the curve propagation method is also faster computationally. Table 3.2 compares the average computation time for the Wang-Fallone and curve propagation method (using the Fast Marching Method and Fast Sweeping Method). As



(a) The original double portal image



(b) Segmentation using the Wang-Fallone method with standard gradient



(c) Segmentation using the Wang-Fallone method with directional gradient and user selected center points



(d) Segmentation using the curve propagation method from initial outline

Figure 3.1: Comparison of segmentation results using each method

Table 3.2: Comparative computation time

Segmentation Method	Image Set	Runtime (in seconds)		
		Easy	Medium	Complex
Wang-Fallone	16-bit images	10.8	10.9	13.3
Fast Marching Method	8-bit images		1.0	
Fast Marching Method	16-bit images		2.3	
Fast Sweeping Method	8-bit images		1.7	
Fast Sweeping Method	16-bit images		4.0	

described above, the Wang-Fallone method could only be reliably run on the 16-bit images. The speed of the Wang-Fallone method has proven to be affected by the complexity of the image so averages for each complexity category are shown. The speed of the curve propagation method is consistent across complexities, but not for the different bit-rates so separate averages are shown.

3.3 Future Work

The current stopping condition for the curve propagation method is to run for a given number of iterations and then stop. Different stopping conditions could be developed such as calculating an overall energy from the curve and stopping when that energy reaches a given level. Another option is to stop when the curve slows down and its movement is very small in the current iteration.

The weaknesses of the level set method are the complexity of the algorithm, which makes it quite slow in comparison to other methods, and the difficulty in finding the right set of parameters for a particular task. As mentioned above a good set of values for the parameters has not been found for double portal image segmentation. To remedy this problem optimization techniques could be used to find such a set of values. The method will still be much slower than the Fast Marching Method and Fast Sweeping Method but with optimal force values it could be used as a final step to enhance the segmentation found by the curve propagation method.

Finally, the curve propagation method should be tested on a much larger set of images to make sure the method is robust enough to be used during a real cancer treatment. If this proves to be the case the method can then be put into use.

Bibliography

- [1] *Canadian Cancer Statistics 2007*, Canadian Cancer Society/National Cancer Institute of Canada, Toronto, Canada, April 2007.
- [2] J. BIJHOLD, K. G. A. GILHUIJS, M. VAN HERK, AND H. MEERTENS, *Radiation field edge detection in portal images*, *Physics in Medicine and Biology*, 36 (1991), pp. 1705–1710.
- [3] S. S. C. BURNETT, G. STARKSCHALL, C. W. STEVENS, AND Z. LIAO, *A deformable-model approach to semi-automatic segmentation of ct images demonstrated by application to the spinal canal*, *Medical Physics*, 31 (2004), pp. 251–263.
- [4] H. DELINGETTE AND J. MONTAGNAT, *Shape and topology constraints on parametric active contours*, *Computer Vision and Image Understanding*, 83 (2001), pp. 140–171.
- [5] K. G. A. GILHUIJS AND M. V. HERK, *Automatic on-line inspection of patient setup in radiation therapy using digital portal images*, *Medical Physics*, 20 (1993), pp. 667–677.
- [6] R. C. GONZALEZ AND R. E. WOODS, *Digital Image Processing*, Prentice Hall, 2nd ed., 2002.
- [7] V. GRAU, A. J. U. MEWES, M. A. RAYA, R. KIKINIS, AND S. K. WARFIELD, *Improved watershed transform for medical image segmentation using prior information.*, *IEEE Trans. Med. Imaging*, 23 (2004), pp. 447–458.
- [8] K. HUMNABADKAR, S. SINGH, D. GHOSH, AND P. BORA, *Unsupervised active contour model for biological image segmentation and analysis*, vol. 2, Oct 2003, pp. 538–542.
- [9] M. KASS, A. WITKIN, AND D. TERZOPOLOUS, *Snakes: Active contour models*, *International Journal of Computer Vision*, 1 (1988), pp. 321–331.
- [10] C. LI, J. LIU, AND M. D. FOX, *Segmentation of edge preserving gradient vector flow: An approach toward automatically initializing and splitting of snakes*, in *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1*, Washington, DC, USA, 2005, IEEE Computer Society, pp. 162–167.
- [11] K. LI, X. WU, D. Z. CHEN, AND M. SONKA, *Optimal surface segmentation in volumetric images-A graph-theoretic approach*, *IEEE Trans. Pattern Anal. Mach. Intell.*, 28 (2006), pp. 119–134.

- [12] X. LI AND G. HAMARNEH, *Modeling prior shape and appearance knowledge in watershed segmentat*, CRV, 00 (2005), pp. 27–33.
- [13] Q. LIN, *Enhancement, extraction, and visualization of 3d volume data*, master's thesis, Linköping University (Sweden), 2003.
- [14] R. MALLADI AND J. A. SETHIAN, *An $O(N \log N)$ algorithm for shape modeling*, Proc. Nat. Acad. Sci. U.S.A., 93 (1996), pp. 9389–9392.
- [15] O. PETRASCU, A. BEL, N. LINTHOUT, D. VERELLEN, G. SOETE, AND G. STORME, *Automatic on-line electronic portal image analysis with a wavelet-based edge detector*, Medical Physics, 27 (2000), pp. 321–329.
- [16] G. SAPIRO, *Vector (self) snakes: A geometric framework for color, texture and multiscale image segmentation*, in International Conference on Image Processing, 1996, pp. 817–820.
- [17] J. A. SETHIAN, *Fast marching methods*, SIAM Review, 41 (1999), pp. 199–235.
- [18] —, *Level set methods and fast marching methods*, vol. 3 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, second ed., 1999. Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science.
- [19] A. TSAI, A. J. YEZZI, W. M. W. III, C. TEMPANY, D. TUCKER, A. FAN, W. E. L. GRIMSON, AND A. S. WILLSKY, *A shape-based approach to the segmentation of medical imagery using level sets*, IEEE Trans. Med. Imaging, 22 (2003), pp. 137–154.
- [20] R. TSAI, L.-T. CHENG, S. OSHER, AND H.-K. ZHAO, *Fast sweeping algorithms for a class of Hamilton-Jacobi equations*, SIAM Journal on Numerical Analysis, 41 (2003), pp. 673–694.
- [21] R. TSAI AND S. OSHER, *Level set methods and their applications in image science*, Communications in Mathematical Sciences, 1 (2003), pp. 623–656.
- [22] H. WANG AND B. G. FALLONE, *A robust morphological algorithm for automatic radiation field extraction and correlation of portal images*, Medical Physics, 21 (1994), pp. 237–244.
- [23] —, *A mathematical model of radiation field edge localization*, Medical Physics, 22 (1995), pp. 1107–1110.