



City Research Online

City St George's, University of London

Citation: Alvino, C. V., Unal, G. B., Slabaugh, G. G., Peny, B. & Fang, T. (2007). Efficient segmentation based on Eikonal and diffusion equations. *International Journal of Computer Mathematics*, 84(9), pp. 1309-1324. doi: 10.1080/00207160701324249

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/4410/>

Link to published version: <https://doi.org/10.1080/00207160701324249>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Efficient Segmentation Based on Eikonal and Diffusion Equations

Christopher Alvino^{*†}, Gozde Unal[†], Greg Slabaugh[†], Bertrand Peny[†], Tong Fang[†]

(Received 00 Month 200x; revised 00 Month 200x; in final form 00 Month 200x)

Segmentation of regions of interest in an image has important applications in medical image analysis, particularly in computer aided diagnosis. Segmentation can enable further quantitative analysis of anatomical structures. We present efficient image segmentation schemes based on the solution of distinct partial differential equations (PDEs). For each known image region, a PDE is solved, the solution of which locally represents the weighted distance from a region known to have a certain segmentation label. To achieve this goal, we propose the use of two separate PDEs, the Eikonal equation and a diffusion equation. In each method, the segmentation labels are obtained by a competition criterion between the solutions to the PDEs corresponding to each region. We discuss how each method applies the concept of information propagation from the labeled image regions to the unknown image regions. Experimental results are presented on magnetic resonance (MR), computed tomography (CT), and ultrasound images and for both two-region and multi-region segmentation problems. These results demonstrate the high level of efficiency as well as the accuracy of the proposed methods.

1. Introduction

Content extraction from images typically relies on segmentation, i.e., extraction of the borders of target structures. Automated segmentation by computer algorithms has been a focus of decades of research [1–3] and remains an active problem in the computer vision literature [4–6]. In practice, the accuracy of segmentation algorithms can be hampered by noise in the image acquisition and the complexity of the arrangement of target objects with respect to their surroundings within the image. In order to achieve robustness to such hindrances, many algorithms demand an increase in computational cost. However, practically useful segmentation techniques should be accurate and computationally efficient for clinical interpretation and so that extensive quantitative analysis can be automated. In this study, highly efficient and mathematically principled techniques are presented to segment the boundaries of closed structures. The techniques are based on ideas of anisotropic information propagation apparent in certain types of partial differential equations (PDEs). This work is motivated by anatomical structures such as lymph nodes, as shown in Fig. 1, whose extraction from medical images, such as magnetic resonance (MR) images, is an important task for subsequent quantitative analysis.

Segmentation methods based on information propagation have been performed using distance functions. For example, in [7], simultaneous propagations are performed to estimate two potentials between two points in order to extract the path made by a vessel. The minimal paths between two points, p_0 and p_1 , are computed by simultaneous propagations from the two points until they meet at a common point p_2 , and by back-propagating to the original two points. They also described an approach to build a path given only a starting point and a given path length. While this approach is suitable for the extraction of tubular structures, our goal is different. Although we also make use of two distance maps, we do not need to extract a minimal path from the point where the two fronts meet, but we seek the result of the competition of the two fronts in reaching a given point. Similarly, in [8, 9] a fast marching algorithm was used for segmenting tubular structures like vessels. A multiphase fast marching algorithm was used in [10] in a Bayesian framework, where all distinct regions are propagated simultaneously according to different velocities, which each depended on the posterior distributions for each region.

There are also similarities between distance function based algorithms and the watershed algorithm. The Eikonal PDE has been used in [11] for modelling watershed segmentation that is constructed from the watershed of the gradient image. Different segmentation results were obtained by changing the flooding criterion [12]. A form of diffusion has been used for image segmentation in [5] by a random walk concept motivated by electric potentials. This technique differs from our approach in that it was introduced in a graph theoretic framework, as has become

^{*}Corresponding author. Email: christopher.alvino@siemens.com

[†]Intelligent Vision and Reasoning, Siemens Corporate Research, Princeton, NJ 08540, USA

popular recently [3, 4], and formulated as system of linear equations solved through conjugate gradient.

Distance functions are intimately linked to level set methods that provide implicit ways to represent boundaries in a way that is free of parameterization and allows for natural topology change. The original applications of level set methods to image segmentation were introduced by Caselles et al. [13], Malladi et al. [14], and Kichenassamy et al. [15]. Distance functions are used in these techniques to enforce a mathematically well-behaved embedding function for the level set of interest, typically the zero level set. Level set methods were soon applied to implementations of the Mumford-Shah functional [16–18].

In this paper we present four methods. The first three methods compute distance functions treating image locations containing edges or higher gradient magnitude as locally slower to propagate information or as having higher local distance. These three methods employ the Eikonal equation and thus can be computed in $O(N \log N)$, where N is the number of image pixels, by the fast marching algorithm [19]. Inspired by the same information propagation concepts, we also present a fourth method based on diffusion PDEs, in which edge information is propagated from the interior of the desired anatomical structure or from the boundary of the region of interest.




Figure 1. Example of magnetic resonance (MR) image with a region of interest (ROI) around a lymph node.

2. Segmentation by Distance Function Competition

We will explain the technique for the case of two-region segmentations and later explain the natural generalization to the case of multiple region segmentations. The first step in the proposed segmentation technique is to compute a distance function for each known image region. Each distance function represents the distance to the nearest of a set of prespecified points interior to the desired structure. The second distance function represents the distance to a set of prespecified points exterior to the structure. We will defer choice of the prespecified interior and exterior points until later, but for now we will state that they should, respectively, be clearly inside or outside the boundaries of the desired structure to be segmented. The local travel cost for each distance function depends on the local image intensity variation. Regions that are more likely to be edges should be interpreted as regions that have higher local distance. After the computation of the two distance functions, a simple competition criterion between the two distance functions determines which image pixels belong to the interior region and which belong to the exterior region.

This concept will be implemented in several different ways. In the first, we weight the distance function directly on the binary map resulting from an edge detection on the image, for instance using a Canny edge detector [20]. In this method, edges correspond to impassable obstacles and the distance function is computed accordingly. The second method generalizes the first method, by defining the local distance as a function of the gradient magnitude of the image. The third method combines the different weights on the distance function. The fourth method is inspired by distance propagation ideas but uses diffusion PDEs as will be explained. The next sections briefly describe the techniques in more detail.

2.1. Eikonal Equation Method

The Eikonal equation,

$$\|\nabla D\| = F, D = 0 \text{ on } G \quad (1)$$

is a well known PDE whose solution, $D : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}$, where n is the dimensionality of the image, represents the arrival time of a moving front with spatially varying speed, $1/F$, that starts at a given set of points, $G \subset \Omega$, at time 0. Here ∇ denote the gradient operator. When the speed of the front is uniform within the domain, the arrival time is proportional to the minimum distance to G , the set of starting points. It is for this reason that the solution to the Eikonal equation is often called a “distance” function. Thus, a common alternate interpretation of the solution to Eq. (1) is that D represents the weighted distance to the set G with locally varying travel cost, F . In this paper, we will use this interpretation and therefore refer to D as a distance function. In our segmentation method the local distance weight will vary accordingly with the presence of image edge or with local intensity variation.

The fast marching algorithm was introduced to yield an efficient solution to the Eikonal equation on a uniform discrete grid [19]. While the theory behind the proposed method holds for continuous image domains with differentiable images, we will herein refer to discrete grid locations and thus, we will use finite difference approximations to the derivatives.

The proposed Eikonal PDE-based methods proceed as follows:

- (i) Initialize for the computation of two distance functions, D^i and D^e , corresponding to the distance from the interior and exterior regions, respectively.
 - a) D^i will be solved on the image domain by setting G to a set of points inside the structure to be segmented. Discretely, this will be done by setting the corresponding pixels to a value of 0 and by labelling the corresponding pixels as *Known*.
 - b) D^e will be solved on the image domain by setting G to a set of points clearly outside the structure to be segmented. Discretely, this will be done by setting the corresponding pixels to 0 and by labelling the corresponding pixels as *Known*.
- (ii) Compute the two distance functions, D^i and D^e , by solving two Eikonal PDEs.

The Eikonal PDE (solved through fast marching): Label the pixels that are neighbors of the already *Known* points as *Trial* pixels. All other image pixels are labeled as *Far* points. Then, until no *Trial* pixels remain, do the following action: take the *Trial* pixel with the lowest distance value, q , label it as a *Known* pixel, and verify that each neighbor pixel to q that is not *Known* is labelled as a *Trial* pixel while updating its value according to the chosen distance function. See [19] for more details.

 - a) Interior: Compute the distance function to the interior set with local travel cost, F , as will be explained in the next sections. The value of each pixel then corresponds to the weighted distance to the interior set and is denoted as D^i . This step is initialized with interior points as *Known* set.
 - b) Exterior: Compute the distance function to the exterior set with local travel cost, F . The value of each pixel then corresponds to the weighted distance to the exterior set and is denoted as D^e . This step is initialized with exterior points as *Known* set.
- (iii) The interior region is considered the set of points where the interior distance is less than the exterior distance, i.e., the interior set is $\{(x, y) : D^i(x, y) < D^e(x, y)\}$ in the case of a two-dimensional image.

The local travel cost, F , of the distance functions are explained in the following sections. We are proposing three different techniques for assigning this travel cost based on the image data. The first technique is based on the presence or absence of edge in the image. The second is based on local intensity variation and the third is a hybrid of the first two techniques.

2.1.1. Fast Marching with Edge Map. Our first approach is to compute the distance function in a way such that edge pixels represent points where the moving front cannot propagate at all. The Eikonal equation is then,

$$\|\nabla D\| = \frac{1}{1 - E}, D = 0 \text{ on } G, \quad (2)$$

where E is the edge map that assumes the value of $1 - \epsilon$ where there are edges and the value of 0 at all other pixels. Note that we are interested in the nature of the solution where $\epsilon > 0$ approaches 0 in order to represent locally infinite travel cost¹. The edge map can be derived from any edge detection algorithm that has binary output. In our results, we use a Canny edge detector [20].

At this point it is important to note how the proposed method differs from using the Canny edge detection algorithm alone. The Canny edge detection algorithm simply reports an edge map that has no *a priori* known topology, i.e., it does not necessarily partition the region of interest into clear interior and exterior regions. The problem of obtaining a labelling for each pixel as an interior or exterior region is thus not solved by edge detection alone. It is for this reason that we propose the use of the competition algorithm.

In the fast marching algorithm the edge pixels are marked as having infinite local travel cost and their initial label is set to *Known*. In this way they will not be processed during the distance function computation. The first column in Figure 2 depicts the two distance functions computed by starting from both the exterior and the interior seed points. The distance is represented in gray scale with white corresponding to high distance and black corresponding to low distance. Note how the exterior distance function has high distance inside the lymph node and low distance outside the lymph node, and the opposite is true for the interior distance function.

2.1.2. Fast Marching with Gradient. In the second method, we treat regions with high gradient magnitude as having high local travel cost, and regions with low gradient magnitude as having low local distance. The Eikonal equation then takes the form:

$$\|\nabla D\| = \|\nabla I\|, D = 0 \text{ on } G, \quad (3)$$

Note that this method has the potential to be more robust to errors in the edge map since it allows moderate levels of intensity variation to affect the local travel cost by a moderate amount instead of necessarily being classified as either edge or non-edge, neither of which classification is completely appropriate. Contrast this to the edge map method in Section 2.1.1 in which pixels that are erroneously considered as edge or non-edge will have a definite negative impact on the distance functions and, as a result, on the final segmentations.

Furthermore, note that by considering the Eikonal equation, $\|\nabla D\| = f(\|\nabla I\|)$, we further generalize this method. Note that f should typically be non-negative and monotonically increasing. Certain choices of the function, f , in particular those that resemble thresholding functions such as the sigmoid, show the relationship between the current method and the method in Section 2.1.1 since edge maps typically resemble such functions of the image gradient magnitude, $\|\nabla I\|$.

The second column in Figure 2 depicts the two distance functions computed in this way.

2.1.3. Combined Method. In the second method explained in Section 2.1.2, which uses the gradient magnitude in the computation of the distance function, there were cases where the distance function was allowed to propagate too quickly through breaks in the boundary of the structure to be segmented. To prevent such leaks and to increase robustness to noise, one can combine the first two methods in Section 2.1.1 and 2.1.2. This corresponds to weighting the distance function by edge information. The method consists of the computation of the edge map, E , to result in a binary image assuming the value $E = 1$ on edge pixels and $E = 0$ elsewhere. This binary image is then directly added to the gradient image by a factor α . The Eikonal equation then takes the form:

$$\|\nabla D\| = (\|\nabla I\| + \alpha E). \quad (4)$$

Specific choice of the parameter, α , depends on the level of trust that can be placed in the edge map, with higher trust corresponding to higher values of α . This will result in increased gradient effects where there are edges as compared with the method in Section 2.1.2.

¹The implementation of this technique does not require the use of limits or a specific choice of ϵ since we can implement the locally infinite travel cost present at edges by simply not allowing information to propagate with the fast marching algorithm.




Figure 2. Top row: Exterior distance function; Bottom row: Interior distance function. Columns 1. edge map method; 2. image gradient magnitude method; 3. diffusion method.

Note that this technique, which assigns additional travel cost to areas where there are definite image edges, is in the spirit of considering the Eikonal equation, $\|\nabla D\| = f(\|\nabla I\|)$, i.e., with non-linear functions of the gradient image. For instance, choosing the function,

$$f(z) = \frac{1}{1 + e^{(-z+t)}} \quad (5)$$

where t represents an image gradient magnitude threshold and can be chosen automatically based on the range values that magnitude of the image gradient assumes. Many choices can be made for the function, f , but a full exploration of these choices is beyond the scope of this paper. We find that the Eikonal equation in Eq. 4 achieves suitable results and combines the previous two methods in a natural fashion.

Part of the flexibility of the proposed methods in this paper are that different PDEs can be used for each region. That is, the interior distance function can be obtained by solving a different PDE than that producing the exterior distance function. This flexibility may, for example, assist in the segmentation of interior regions that are textured. One can add, to the interior distance function, some interior intensity based term, which will smooth the local gradient and decrease some texture or noise influence. Due to the nature of the main application, i.e., lymph node segmentation, we do not smooth the exterior distance function in a similar fashion since exterior regions may include other structures that may interfere with the segmentation. In order to achieve this, we compute the mean intensity of a set of points adjacent to the foreground seed points as $\hat{I}$. The image at each pixel, p , will then have a local weight of $(I(p) - \hat{I})^2$, which we add to the local travel cost in the Eikonal equation for the interior region with a weighting parameter, β as follows,

$$\|\nabla D^i\| = \left(\|\nabla I\| + \alpha E + \beta (I - \hat{I})^2 \right), \quad (6)$$

where again, E is the binary edge map and $\hat{I}$ is the mean intensity of points adjacent to the interior region seed points. Thus, the final combined method assumes the use of Eq. 4 for computation of the exterior distance function and Eq. 6 for the computation of the interior distance function.

2.2. Diffusion Equation

The technique proposed in this section differs from the above techniques in that we use a diffusion equation to propagate image information rather than the Eikonal equation. Note that although the nature of information propagation in diffusion equations is significantly different than that in the Eikonal equation, particularly that diffusion equations propagate information with infinite speed, we find that this technique propagates information in a similar

manner to the previously mentioned distance function based techniques. Perhaps the reason for this similarity is that diffusion equations propagate information in a gradual manner that depends on the proximity to the information source. Diffusion equations propagate information with infinite speed, but in a way such that the majority of the effect is local.

The linear heat equation on a function $D(\bar{x}, t) : \Omega \times [0, \infty) \rightarrow \mathbb{R}$, where $\bar{x} \in \Omega$, is given by $\frac{dD}{dt} = \Delta D$. Here, Δ denotes the Laplacian operator. We consider initial conditions $D(\bar{x})|_{t=0} = D_0(\bar{x}) = 0$ and Dirichlet boundary conditions $D = 1$ on G for a prespecified seed set $G \subset \Omega$. For more information on the linear heat equation, initial conditions, and Dirichlet boundary conditions, we refer the reader to [21]. An update equation corresponding to the finite difference approximation to this equation for two-dimensional images, that is obtained by implementing a forward Euler numerical scheme with the maximally stable time step ($\Delta t = 0.25$) is,

$$D(x, y) \leftarrow D(x, y) + \Delta t (\Delta D(x, y)) \quad (7)$$

$$\begin{aligned} D(x, y) \leftarrow & \frac{1}{4}D(x+1, y) + \frac{1}{4}D(x-1, y) \\ & + \frac{1}{4}D(x, y-1) + \frac{1}{4}D(x, y+1), \end{aligned} \quad (8)$$

hence diffusing edge information from the boundaries towards the non-boundary regions.

Inspired by the Eikonal equation and fast marching techniques, where we propagate the information from the boundaries or the seeds of the image domain towards unlabeled points, diffusion equations can also be used for segmentation by creating two smooth distance functions¹, one for the interior seed points and one the exterior seed points. For the interior distance function, D^i , the boundary conditions are set to 1 at the interior seed points and the function is initialized to a value of 0 at all other points in the image domain. For the exterior distance function, D^e , the boundary conditions are set to 1 at the exterior seed points and the function is initialized to a value of 0 at all others points in the image domain.

To introduce image dependent terms to the diffusion equation, we propose the use of an anisotropic diffusion that depends on the local image variation, allowing less diffusion in directions where the image derivative is lower and more diffusion where the image derivative is higher. The definition of the four one-sided image derivatives around a pixel are given by

$$\begin{aligned} I_x^-(x, y) &= I(x, y) - I(x-1, y), \quad I_x^+(x, y) = I(x+1, y) - I(x, y) \\ I_y^-(x, y) &= I(x, y) - I(x, y-1), \quad I_y^+(x, y) = I(x, y+1) - I(x, y) \end{aligned}$$

We can create an image-based discrete diffusion equation by introducing the image-driven weights to the discrete heat equation as follows,

$$\begin{aligned} D(x, y) &= \frac{w^E}{\sum_i w^i} D(x+1, y) + \frac{w^W}{\sum_i w^i} D(x-1, y) \\ &+ \frac{w^N}{\sum_i w^i} D(x, y-1) + \frac{w^S}{\sum_i w^i} D(x, y+1), \end{aligned} \quad (9)$$

$$\begin{aligned} w^E &= e^{-\gamma(I_x^+)^2}, \quad w^W = e^{-\gamma(I_x^-)^2}, \\ w^N &= e^{-\gamma(I_y^-)^2}, \quad w^S = e^{-\gamma(I_y^+)^2}, \quad i \in \{E, W, N, S\}. \end{aligned}$$

Note that γ represents a damping coefficient that affects the level of anisotropy inherent in this method. Higher values of γ allow for greater anisotropy, i.e., allow for the information diffusion to be more sensitive to differences in image intensity across the image. We have found a reasonable range for this parameter is $0.001 < \gamma < 0.01$.

Hence, using the set of seeds for the exterior region and the interior region as two distinct set of boundary conditions, we estimate the two distance functions, D^e and D^i , corresponding to the exterior and interior after a

¹Note that for consistency with the previous sections, we refer to these functions as distance functions, however that they more closely resemble heat functions as is typically the case with such diffusion equations.

set amount of diffusion time. Similar to our approach using Eikonal equation, we form the segmentation map by considering the interior region to be the set of points where the interior distance function is higher than the exterior distance function. The third column in Figure 2 depicts the resulting distance functions estimated by the diffusion method. We run the diffusion for a sufficiently long time for all pixels to be affected by the diffusion but for a short enough time for the diffusion to be practically useful, i.e., to avoid the constant solution $D = 1$ along the entire domain. In practice, we found that $k = 1000$ forward Euler iterations produced suitable results and that the results were not particularly sensitive to moderate variations in k .

This image-weighted diffusion we seek for our distance function D is similar to that of the diffusion equation presented in the work of Perona and Malik [22] who used anisotropic diffusion for filtering images based on the direction of the image gradient. Using a weighted diffusion equation based on image gradients as, $\partial I / \partial t = \nabla \cdot (w(|\nabla I|) \nabla I)$, their purpose is to diffuse intensities of the original image, I , in an edge-preserving manner and not to derive distance functions, D^e and D^i , as we do.

3. Experimental Results

3.1. Two Region Segmentations

The algorithm is not sensitive to the placement of the interior and exterior seed points. It is possible to use any set of exterior or interior seed points as long as they are clearly outside or inside of the target structure, respectively. Of course, a carefully hand labelled contour placed at a uniform distance from the outside of the structure may be ideal for the exterior seed points. However, for simplicity, and to show the robustness of the proposed method to the choice of initial contour, we opted to use a simple mouse drag operation on the image that sets exterior seeds in the form of a 2D rectangular border, as shown in Fig. 1, then the interior seeds are automatically set to the set of pixels in the center of this rectangle. This type of 2D initialization is used in both the 2D and 3D experiments, and is simple and fast for the user since it essentially only requires specifying two points, i.e., the top left point and the bottom right point of the rectangular region of interest.

In Figure 3, sample segmentation results (labeled as blue contours) are presented for different lymph nodes in MR images. In analyzing the results based on the edge map algorithm, we see that in some cases the segmentation is not as precise as the other methods. The Canny edge detector propagates strong edges and discards the weak ones, and this leads to either “holes” in the edge map as in row 1, or edge noise as in rows 3, 4 and 5. This will directly influence the distance functions and in turn, will influence the final segmentation. Still, the result is acceptable and can be used as a suitable fast initialization to a more sophisticated segmentation algorithm. Those errors are reduced by our second approach that uses image gradient in the Eikonal PDE. The distances found are more robust to errors in the edge map functions, and our segmentation matches the desired structure more closely. In cases where a strong edge is situated near the edges of the lymph node, but yet external to the lymph node itself, the method presented in this paper may be slightly attracted to it, such as the bottom right the image in row 1 and the entire surrounding region for the image in row 5. Note that such errors can be fixed by assuming more than two regions are present in the image. We will show example of this in the next section.

The diffusion method performs well when images edges are strong and is robust to higher levels of image noise especially the noise that occurs at a single pixel as is commonly referred to as “salt and pepper” noise. The reason for this is the ability of the diffusion equation to propagate information around a single pixel that, due to noise, has an abnormally high intensity value. This property does not hold for the image gradient based techniques since such noise affects the image gradient in a neighborhood around the pixel, thereby creating a larger region of high local travel cost. However, this method is prone to error when the target objects are merged with other external structures containing high edge content. A startling example of this can be seen in the fifth row of Fig. 3. The high amounts of external intensity variation that are external to the target object will often produce such unsatisfactory results.

Finally, our combined method produces the best results, achieving a suitable tradeoff between the results in the edge map method and the image gradient method, even in difficult nodes. This method specifically prohibits the propagation of information where edge is detected, whereas such information often propagates in the image gradient method. However, when edge is not detected, it allows for varying degrees of information propagation depending on the magnitude of the image gradient. We can see in the results Fig. 3 that a suitable tradeoff between the first two methods is achieved and excellent segmentations are produced even in such noisy and complex images.

3.2. Comparison with Ground Truth

The results are confirmed by the statistics we found during our tests as shown in Table 1. The ground truth of each node segmentation was hand labelled by our own learned interpretation from clinicians. We estimated the empirical probabilities of pixels falsely accepted as foreground (Type I error) or pixels falsely rejected as foreground (Type II error) on the resulting contours of the presented four segmentation methods compared with the manually delineated node contours. These results are compiled over a database of 50 different regions of interest containing lymph nodes in T2 weighted and T2-star weighted magnetic resonance images. The seed regions were kept constant for each node and only the segmentation method varied.

The very low value in the Type I error of the edge map method is explained by its preference to label pixels as foreground. This is highlighted by the large probability of Type II error. On the other hand, the gradient and diffusion methods are more prone to erroneous information propagation around edges and thus have then a higher Type I error. However, it is important to note that the overall probability of error is lower in the gradient and diffusion methods than in the edge map method. Finally, the combined method yields the lowest overall error rate.

3.3. Further Experimental Results

Segmentation in 3D through Eikonal PDEs is easily achieved by extending the fast marching, and the gradient computations to three dimensions. Example results from two nodes are shown in Fig. 4. We also show the effectiveness of the proposed segmentation techniques in other imaging modalities. The example in Fig. 5 on a Computed Tomography (CT) image shows a selected region of interest containing a liver tumor on the left and the segmentation of the tumor as shown on the right.

In Fig. 6, we show segmentation results on the same slice of the same CT volume as in Fig. 5. However, in Fig. 6 we vary the region of interest (ROI) substantially to show the robustness of the segmentation result to different possible ROIs that are likely to have been designated for the tumor in question. In the first row, we show the performance of the algorithm with expected input ROIs, i.e., shifted around the tumor, and the corresponding stable segmentation results. In the second row, we show the performance of the algorithm with extreme, not commonly expected ROIs. In these cases, the proposed algorithm still captures the tumor boundaries in a reasonable way.

The 3D tumor extraction results are shown in Fig. 7. Such 3D segmentations can be useful when trying to determine the volume of a tumor, for example to assess the longitudinal efficacy of treatment. Figure 8 shows an example of a breast mass segmentation in an ultrasound image. As we can see, ultrasound images have speckle noise that hampers segmentation, and therefore we had to pre-process the image with high level of smoothing to reduce it. However, the method remains accurate and efficient when the proper level of smoothing is applied. The results show that our algorithm works for different types of structures that have a clear interior and exterior and may be tuned for applications other than lymph node segmentation.

3.4. Computation Speed

The Eikonal PDE-based approaches presented in this paper, as expected, are very fast due to the fast computation of the fast marching algorithm. On a 60 by 60 pixel region of interest, the segmentation is complete in less than 0.03 seconds for the 2D algorithm, and 0.76 seconds for the 3D algorithm on a 60 by 60 by 60 pixel region of interest. All results are reported with the algorithms running on a Pentium 4 processor operating at 2.4 GHz. With the diffusion PDE, the segmentation is completed in 1.75 seconds for a 2D implementation. Although we extended the diffusion approach to 3D, the computation times increased to approximately 1 to 2 minutes, therefore, we have not used the diffusion-based approach for the 3D experiments. Considering the high level of accuracy of the proposed algorithms in noisy and complex images, and the relative simplicity of the user supplied labelling, these algorithms achieve a very high computational efficiency.

3.5. Multi-Region Segmentation

To illustrate the principled nature of the proposed segmentation techniques, particularly the distance function techniques, we explain the generalization of this technique to segmentation of multiple regions. We will show examples

on regions of interest that have three clear regions. The generalization of this technique is mathematically natural. Instead of choosing only interior and exterior seed regions as in the two region segmentation, the multiple region segmentation technique allows for any number of seed regions where each seed region should correspond to a structure of interest within the image. For each seed region, we then compute a distance function from that seed region as before. The segmentation labels are assigned by determining which distance function has the lowest value for a given pixel. That is, for a each label index, i , the set of points corresponding to that region, L_i is given by,

$$L_i = \{(x, y) : D_i(x, y) \leq D_j(x, y), \forall j \neq i\}. \quad (10)$$

Note that some pixels may be defined with multiple labels with such a definition, however, these pixels are typically pixels that are on the border between two or more regions and can be thus considered border pixels, or can be assigned to one region of the other without any loss in the utility of the method. This generalization is both mathematically natural and matches our geometric intuition.

Figure 9 shows some examples of three region segmentations. Each row in this figure corresponds to a different region of interest in the MR lymph node segmentation application. The left column shows the initial region of interest (ROI) along with the user specified seeds in yellow. Note that there are typically two foreground seeds and a single background seed surrounding the entire region. In the final row however, there are two foreground seeds, one of which is split along the two dark blood vessels, and one of which is in the white lymph node.

The middle column shows the gradient of the smoothed ROI, and the right column shows the corresponding final segmentation result using the method of distance function locally weighted by the image gradient alone, i.e., $\|\nabla D\| = \|\nabla I\|$. For each image, one seed region is the rectangular box surrounding the structure and two interior seed regions are chosen, one inside of each region of interest as shown in the first column. Using a multi-region segmentation method is particularly helpful in that it allows for explicit modelling of structures that have significant edge information and that are external to the structure of interest, rather than having these structures confound the distance functions and thus, create errors in the final two-region segmentation. This is particularly the situation for the lymph node application, where lymph nodes are found in the vicinity of the vessels.

We should finally note that, as shown in the example in Fig. 6, the multi-region version of the algorithm shows similar levels of flexibility when the center points are no longer directly in the center of the object to be segmented. This level of flexibility is attained by the design of the distance functions, which are not very sensitive to small shifts in the center points.

4. Conclusion

In conclusion, we presented efficient image segmentation techniques based on ideas from the Eikonal and diffusion PDEs, by computing the distance functions for the exterior and interior regions, and by determining the final segmentation labels by a competition criterion between the distance functions for reaching a given point. Each method has its pros and cons, according to the image characteristics, but our experiments demonstrated that among the presented methods, the combined edge map and image gradient method achieves the most accurate segmentations, and hence the best utility when compared to the other three methods. We have additionally shown the natural generalization of the two-region method to segmenting multiple regions. The segmentations resulting from this algorithm are both fast and accurate and are beneficial for clinical applications that require segmentation with a minimal amount of user interaction.

Acknowledgements

We thank Dr. M. Harisinghani, Dr. R. Weissleder at Massachusetts General Hospital (MGH) in Boston, Dr. J. Barentsz at University Medical Center in Nijmegen, Netherlands, for clinical motivation, feedback and providing data, and Dr. R. Seethamraju for discussions, Dr. R. Krieg at Siemens Medical Solution for support of this work.

Table 1. Estimate type I and type II error probabilities over a database of 50 nodes

	Edge Map Method	Gradient Method	Diffusion Method	Combined Method
Type I	0.015	0.247	0.256	0.081
Type II	0.453	0.115	0.189	0.257

References

- [1] Geman, S., Geman, D.: Stochastics relaxation, Gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **6** (1984) 721–741
- [2] Blake, A., Zisserman, A.: *Visual Reconstruction*. MIT Press (1987)
- [3] Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **22**(8) (2000) 888–905
- [4] Boykov, Y., Jolly, M.: Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In: *ICCV*. Volume 1. (2001) 105–112
- [5] Grady, L., Lea, G.: Multi-label image segmentation for medical applications based on graph-theoretic electric potentials. In: *ECCV, Workshop on MIA and MMBIA*. (2004)
- [6] Grady, L., Schwartz, E.L.: Isoperimetric graph partitioning for image segmentation. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **28**(3) (2006) 469–475
- [7] Deschamps, T., Cohen, L.D.: Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *MIA* **5**(4) (2001) 281–299
- [8] Deschamps, T., Cohen, L.D.: Fast extraction of tubular and tree 3D surfaces with front propagation methods. *16th International Conference on Pattern Recognition* (2002)
- [9] Cohen, L.D., Kimmel, R.: Global minimum for active contour models: A minimal path approach. *Int. J. Computer Vision* **24**(1) (1997) 57–78
- [10] Sifakis, E., Garcia, C., Tziritas, G.: Bayesian level sets for image segmentation. *J. Vis. Commun. Im. Repres.* **13**(1) (2002) 44–64
- [11] Meyer, F., Maragos, P.: Multiscale morphological segmentations based on watershed, flooding, and Eikonal PDE. *Proc. Scale Space Theories in Computer Vision* (1999) 351–362
- [12] Sofou, A., Maragos, P.: PDE-based modeling of image segmentation using volumic flooding. *IEEE Int. Conf. on Image Processing* (2003) 431–434
- [13] Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. In: *Proc. IEEE Int. Conf. Computer Vision*. (1995) 694–699
- [14] Malladi, R., Sethian, J.A., Vermuri, B.C.: Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **17**(2) (1995) 158–175
- [15] Kichenassamy, S., Kumar, A., Olver, P.J., Tannenbaum, A., Yezzi, A.J.: Gradient flows and geometric active contour models. In: *IEEE Int. Conf. Computer Vision*. (1995) 810–815
- [16] Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Trans. Image Process.* **10**(2) (2001) 266–277
- [17] Tsai, A., Yezzi, A., Willsky, A.: Curve evolution implementation of the Mumford-Shah function for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. Image Process.* **10**(8) (2001) 1169–1186
- [18] Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Applied Math* **42** (1989) 577–685
- [19] Sethian, J.: *Level Set Methods and Fast Marching Methods*. Cambridge Univ. Press (1999)
- [20] Canny, J.: A computational approach to edge detection. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **8**(6) (1986) 679–698
- [21] Evans, L.C.: *Partial Differential Equations*. American Mathematical Society (2000)
- [22] Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis, and Machine Intelligence* **12**(7) (1990) 629–639




Figure 3. Segmentation Results. Columns(a–f): a. ROI image; b. Node manually delineated; c. Edge Map Method; d. Gradient Method; e. Diffusion Method, f. Combined Method.




Figure 4. 3D Segmentation of anatomic structures based on Eikonal PDEs.




Figure 5. A liver tumor is segmented using the Combined Algorithm on a CT volume.




Figure 6. Liver tumor from same CT volume as in Fig. 5 segmented with different regions of interest (ROIs) of varying size and shape. The resulting segmentations of the liver tumor are shown. First row: with expected input ROIs, shifted around the tumor, and stable segmentation results. Second row: even with extreme, not commonly expected ROIs, the proposed algorithm still captures the tumor boundaries in a reasonable way.




Figure 7. 3D Segmentation results on CT sequences of Fig. 5




Figure 8. A breast mass segmented using the Combined Algorithm on a Ultrasound image.




Figure 9. Example three-region segmentations with lymph nodes and blood vessels as the two foreground regions respectively.