Unseeded region growing for 3D image segmentation

Zheng Lin[1] **Jesse Jin**[1] **Hugues Talbot**[2]

[1] School of Computer Science & Engineering
University of New South Wales, Sydney 2052
zlin{jesse}@cse.unsw.edu.au

[2] CSIRO Mathematical and Information Sciences
Locked Bag 17, North Ryde 1670, New South Wales
Hugues.Talbot@cmis.csiro.au

Abstract

Unseeded region growing is a versatile and fully automatic segmentation technique suitable for multispectral and 3D images. This approach integrates region-based segmentation with image processing techniques based on adaptive anisotropic diffusion filters.

The segmentation method is fast, reliable and free of tuning parameters. It is indeed a general purpose segmentation method and has been successfully applied in a range of image analysis tasks. This paper describes the algorithm, and briefly discusses its properties and applications. Segmentation results will also be shown at the end of the paper.

1 Introduction

Segmentation is the process in which an image is divided into constituent objects or parts. It is often the first and most vital step in an image analysis task. Effective segmentation can usually dictate eventual success of the analysis. For this reason, many segmentation techniques have been developed by researchers worldwide [6].

Segmentation of intensity images usually involve four main approaches, namely thresholding, boundary detection, region-based and hybrid methods.

Thresholding techniques [12] are based on the postulate that all pixel whose value lie within a certain range belongs to one class. Such methods neglect all of the spatial information of the image and do not cope well with noise or blurring at boundaries.

Boundary-based methods are sometimes called edge-detection [4], because they assume that pixel values change rapidly at the boundary between two regions. The basic method is to apply a gradient filter to the image. High values of this filter provide candidates for region boundaries, which must then be modified to produce closed curves representing the boundaries between regions.

Region-based segmentation algorithms postulate that neighboring pixels within the same region have similar intensity values, of which the split-and-merge [7] technique is probably the most well known. The general procedure is to compare a pixel with its immediate surrounding neighbors. If a criterion of homogeneity is satisfied, the pixel can be classified into the same class as one or more of its neighbors. The choice of homogeneity criterion is therefore critical to the success of the segmentation.

Hybrid methods combines one or more of the abovementioned criteria. This class includes the morphological watershed [9] segmentation, variable-order surface fitting [3] and active contour [8] methods. The watershed method is generally applied to the gradient of the image. The gradient image can be considered as a topography with boundaries between regions as ridges. Unlike the boundary-based methods, the watershed is guaranteed to produce closed boundaries even if the transition between regions are of variable strength or sharpness. However, it encounters difficulties with images in which regions are both noisy and have blurred or indistinct boundaries. The variable-order surface fitting method starts with a coarse segmentation of the image into several primitives, which are refined by iterative region growing procedure. Active contour models are based on gradient information along the boundary between regions, but they are useful only when a good initial estimate is present.

We present a new segmentation method known as "unseeded region growing" (URG) which is based on conventional region-growing postulate of pixel similarities within regions. However, unlike its conventional counterpart, it does not rely on fine-tuning homogeneity parameters, nor does it require manual inputs known as seeds. The result is an algorithm which is robust, easy to use and can readily incorporate higher level knowledge of the image composition through the choice of region threshold, which can be conceptualized as the contrast between different regions.

Despite this significant improvement in image segmentation technique, the segmentation result can still be limited by low signal to noise ratio (SNR) or contrast to noise ratio (CNR).

Filtering techniques can often be employed to reduce the amplitude of noise fluctuations, but conventional linear filters such as low pass filters will usually degrade sharp details such as lines or edges. The filtering also does not preserve region boundaries or small structures, thus the resulting images appears blurry. These undesirable effects can be reduced by using nonlinear filters, the most common one being the median filtering. However, median filtering results in a loss of resolution due to suppression of fine details.

Image degradation by blurring or by artifacts due to filtering are not desirable. An ideal filter should enhance morphological definition of the image by sharpening discontinuities and remove noise in homogeneous regions while preserving object boundaries and fine details.

Recent developments based on anisotropic filtering [11][5] satisfies most of the above criterion. This filtering technique overcomes the drawbacks of conventional spatial filtering techniques and significantly improves image quality. Extensions of this method are especially appropriate for enhancement of various types of multi-channel and 3D image data.

Anisotropic filters can be used as a preprocessing step to effectively eliminate noise in the input image as much as possible, so that the effects of noise in the region growing procedure can be minimized. Combining anisotropic filtering with unseeded region growing results in a powerful segmentation routine which is fast, reliable and consistent for a wide range of image analysis tasks.

2 Related Work

Seeded region growing (SRG) [1] is a well known region-based segmentation method that segments intensity images into regions based on a marker set (seeds). The selection of seeds determines what is a feature of interest in the image and what is irrelevant. As the name implies, region growing is a procedure that groups pixels into larger regions. Border pixels are added to regions in an order that depends on the similarity between the pixel and the marked region, thus the segmentation result is highly dependent on the choice of seeds.

The SRG algorithm always segments an image into regions with the property that each connected component of a region contains exactly one of the seed classes, and the regions are chosen to be as homogeneous as possible subject to this constraint. This methodology can be extended to multi-dimensional and multi-spectral images as well.

SRG starts with a set of seed points, and regions are grown from these seeds by appending to each seeded region those neighboring pixels whose properties are most similar to the region. The homogeneity criterion is the minimum difference between the graylevel of a pixel and the average graylevel of its assigned region. Once the criterion is met, the pixel is said to belong to the same region as one or more of its neighbors. The segmentation is iterative: at every step, the most suitable pixel from the border set is selected according to the homogeneity criterion, and its neighboring points are added to the border set. This process is repeated until all pixels have been allocated to one of the regions.

2.1 Known Issues

There are a number of known issues associated with SRG scheme, some of these have been pointed out in [2].

- A number of implementations of SRG only approximate the behavior of the original algorithm. This can result in scan-order dependencies and can have significant impacts on small regions.
- A good segmentation result depends on a set of "correct" choice for the seeds. When the input images are noisy, the seeds may fall on atypical pixels that are not representative of the region statistics. This can lead to erroneous segmentation results.
- The seed selection process in itself requires manual interventions, and is error-prone. Even though automatic segmentation can be achieved in a limited sense, it is application specific and will require domain-specific knowledge and training sets.

All of these make SRG unsuitable for vision-type segmentation where a priori knowledge is limited.

3 Unseeded Region Growing

We are proposing a region-based segmentation technique similar to SRG, except that no explicit seed selection is necessary: the seeds can be generated by the segmentation procedure automatically. Therefore, this method can achieve fully automatic segmentation with the added benefit of robustness from being a region-based segmentation.

Formally, the segmentation process initializes with region A_1 containing a single image pixel, and the running state of the segmentation process consist of a set of identified regions, A_1, A_2, $\cdots$, A_n. Let T be the set of all unallocated pixels which borders at least one of these regions:

$$T = \left\{ x \notin \bigcup_{i=1}^{n} A_i \wedge \exists k : N(x) \cap A_k \neq \oslash \right\}$$

where $N(x)$ are immediate neighboring pixels of point x.

Further, we define a difference measure

$$\delta(x, A_i) = |g(x) - \text{mean}_{y \in A_i}[g(y)]|$$

where $g(x)$ denotes the image value at point x, and i is an index of the region such that $N(x)$ intersect A_i.

The growing process involves selecting a point $z \in T$ and region A_j where $j \in [1, n]$ such that

$$\delta(z, A_j) = \min_{x \in T, k \in [1,n]} \{\delta(x, A_k)\}$$

If $\delta(z, A_j)$ is less than the predefined threshold t, then the pixel is added to A_j. Otherwise, we must choose the most substantially similar region $\mathcal{A}$ such that

$$\mathcal{A} = \text{argmin}_{A_k} \{\delta(z, A_k)\}$$

If $\delta(z, \mathcal{A}) < t$, we can assign the pixel to $\mathcal{A}$. If neither of these two conditions above apply, then it is apparent that the pixel is significantly different from all the regions found so far, so a new region, A_{n+1} would be identified and initialized with point z. In all three cases, the statistic of the assigned region must be updated once the pixel has been added to the region.

The URG segmentation procedure is inherently iterative, and the above process is repeated until all pixels have been allocated to a region. For convenience, the initial starting point has been chosen to be the first image pixel, but preliminary investigations have suggested that the starting position does not have a significant influence on the segmentation result.

3.1 Implementation

To ensure correct behavior with respect to the homogeneity criterion, the region growing operation requires the determination of the "best" pixel each time a region statistic is changed. This would be an extremely expensive operation if the δ values for all pixels in T are re-evaluated and the priority queue resorted. A more efficient solution has been outlined in [2]. The trick is to arrange the pixels in a structure that permits rapid search for the best pixel.

We used a combination of splay queue and heap structure to keep track of all pixels currently under consideration.

The splay queue structure is essentially a binary tree containing candidate pixels of each region, sorted by the pixel intensity value. Each node of the binary tree stores pixels with the same intensity in a FIFO queue. The search for "best" pixels of the region (ones that are are closest to the region statistic) resorts down to a simple binary search. This idea can be extend to multispectral pixel data.

When selecting the global best pixel, the set of best candidate pixels for each region are found with their δ values calculated. These regional pixels are stored in a global priority queue, implemented as a heap, with priority dependent on δ. The best pixel over all regions is then simply the candidate pixel that has the smallest δ to its assigned region.

Once the global best candidate has been selected, it is removed from the priority queue and added to an appropriate region. The statistic of the corresponding region is then updated, and neighbors of that pixel added to the regional queue. New regional candidate pixel is found and added to the global queue. The above process will repeat until all pixels have been classified.

It is interesting to note that the hierarchy of the data structure used in this algorithm implies smaller queues that only needs to be changed when their associated region changes. The size of the global queue is dependent on the number of regions, which will typically be small.

Scan order dependency can be eliminated by processing all pixels with the same priority in parallel, and updating the region statistics only when a pixel comes out of the queues with a different value.

3.2 Discussions

There are still two issues relating to the URG method, due to the fact that almost no knowledge of the image content is required for the segmentation.

The most important issue for consideration is the threshold, which is vital to the success of the segmentation process. Threshold is closely correlated with the contrast in the image, so selection of the threshold is a process of analyzing contrast. Several automatic threshold selection methods exist, but they usually require bi-peak or multi-peak histograms, which may not be applicable in many cases. Therefore, manual specification of the threshold is necessary for the segmentation procedure until better threshold selection algorithms becomes available.

Another potential issue is the bias effect on growing regions. Although the region growing procedure in itself is not scan order dependent, the classification process tends to be somewhat biased towards the regions that are discovered earlier. During the segmentation process, we only have the statistics of the region discovered so far, therefore we cannot make reliable predictions on the distributions in the areas yet to be classified. In other words, we may not have the optimal classification for a particular pixel in the sense that there may exist a better region assignment for that given pixel, but those pixels belonging to the "desirable" region may still be in the region queues pending for consideration, or they may not even have been encountered at all.

Several adjustments can be made to offset this bias effect. The first method is to use a two-pass scheme that that takes the initial classification as a guidance for the subsequent segmentation. With the knowledge of region location and their statistical distribution, the refined classification should be more accurate. The second method is simpler: when a new region is discovered, its neighboring pixels that have been previously classified will be put into consideration again. If such pixels are better described by the new region, then they will be reclassified. This process will continue until no changes are being made on existing classifications. The third method has a mechanism similar to the second method, and is based on the observation that bias mainly takes effect around the borders of adjacent regions. Therefore, adjustments are only necessary for the border pixels, and can take place after the initial classification is complete, when the region statistics are more stable.

4 Anisotropic Filtering

Anisotropic filtering performs piecewise smoothing of the original image signal. Its strength lies in the fact that it deals with local estimates of the underlying image structures, which are highly flexible. Discontinuities can be preserved and their positions will not be affected.

The filtering process can be formulated mathematically as diffusion, which is suppressed or stopped at boundaries by locally adaptive diffusion strengths, described as follows:

$$\frac{\partial}{\partial t} I(\overline{x}, t) = \operatorname{div}\left(c(\overline{x}, t) \times \nabla I(\overline{x}, t)\right)$$

The function $I(\overline{x}, t)$ is taken as the image intensity. The diffusion strength is controlled by $c(\overline{x}, t)$, where $\overline{x}$ represents the spatial coordinates, and t is the enumeration of iteration steps. The diffusion function $c(\overline{x}, t) = f(|\nabla I(\overline{x}, t)|)$ depends on the magnitude of the gradient of the image intensity, which mainly diffuses within regions and does not affect region boundaries at locations of high gradients. Two diffusion functions have been used frequently:

$$c_1(\overline{x}, t) = \exp\left(-\left(\frac{|\nabla I(\overline{x}, t)|}{\sqrt{2}\kappa}\right)^2\right)$$

$$c_2(\overline{x}, t) = \frac{1}{1 + \left(\frac{|\nabla I(\overline{x}, t)|}{\kappa}\right)^{1+\alpha}} \mid \alpha > 0$$

The parameter κ is chosen according to the noise level and the edge strength. The relationship between the parameter κ and the gradient ∇I can be explained by the flux function $\phi(\nabla I) = c \times \nabla I$. For example, maximum flux for c_1 is generated when the gradient equals to κ. Below κ, the flow reduces to zero, because only minimal flow takes place in homogeneous regions. Above κ, the flow again decreases to zero, stopping diffusion at locations of high gradients. A proper choice of the diffusion parameter not only preserves, but also enhances edges while being numerically stable.

To filter discrete signals, it can be shown that $I(t + \Delta t) \approx I(t) + \Delta t \times \frac{\partial}{\partial t} I$. Local gradient estimates are calculated as differences between neighboring image pixels instead of differentiation. Stability of the iterated processing can be obtained by choosing a proper integration constant Δt.

Dimension	Connectivity	Δt_{max}
1-D	2	$1/3$
2-D	4	$1/5$
	8	$1/7$
3D	6	$1/7$
	18	$1/13$
	26	$3/47$

Table 1: Integration constant Δt for different neighborhood structures

4.1 Convergence

The diffusion process described above did not incorporate a convergence criterion, so it does not provide a mechanism to inform when the diffusion process should stop, or say after how many iterations we can obtain the best result. It can even be argued that the image would converge to a constant at the limit of infinite time.

Nordstrm [10] proposed an algorithm that unifies the concepts of anisotropic diffusion and variation regularization. It only differs from anisotropic filtering by an additional term expressing the deviation between the original $I(t_0)$ and the filtered images. The discrete formulation can be described as $I(t + \Delta t) \approx I(t) + \Delta t \times \left(\frac{\partial}{\partial t} I + (I(t_0) - I(t))\right)$.

The bias term $I(t_0) - I(t)$ is responsible for the mathematically sound convergence property, but it does not influence the local decision of enhancement or blurring. No number of iterations has to be pre-specified, since the images can be repeatedly filtered until a steady state is reached.

4.2 Adaptive Filtering

Since an image usually consist of several objects with different contrasts, it is important to be adaptive to different areas with different gradients. The general idea behind adaptive filtering is to apply a versatile operator which adapts itself to the local topography of the image being smoothed.

The diffusion process depends critically on the value of κ, as explained previously. The value of κ, in essence, has the same character of contrast. Therefore, we can calculate the gradient in different areas of the image and choose κ in such a way that it is less than the edge gradient and larger than the average value of noise gradient. The first requirement prevents the edge from being blurred and the second ensures the removal of noise.

The algorithm for estimating κ can be formulated as follows

$$\kappa = \max\left(\frac{\sum \nabla_l I}{N_l}\right)$$

where l denotes pixel value and N_l is the number of pixels with value l, so $\sum \nabla_l I / N_l$ is the average gradient of those pixels with value equal to l. In other words, we choose the maximum average gradient as the κ value, assuming that the gradient at edges is larger than that of other areas if noises can be neglected.

Several measures can be taken to reduce the influence of noise gradients on the estimation [13], including isolated noise suppression and regression methods.

4.3 Extensions

The diffusion process can be easily extended to 3D and multispectral data.

Three dimensional diffusion follows directly from the original anisotropic diffusion equation, where $\overline{x}$ corresponds to (x, y, z). The total flow contribution at each node is taken from within a $3 \times 3 \times 3$ neighborhood window. The increased number of samples results in a much better noise reduction and enhancement of 3D discontinuities, which allows for more accurate preservation of 3D structures. When the voxel structure is non-cubic, the larger distance in one dimension must be taken into account by setting the correct distance into calculation for the gradients and the diffusion coefficients.

When dealing with multichannel (or multispectral) data, each voxel location can represent a vector of information, with each component representing different physical properties. Filtering can be performed simultaneously across all channels, assuming perfect spatial coincidence. Channel coupling is achieved by combining the corresponding diffusion coefficients, which inherently preserves correlating and contrasting effects of image intensities among those channels.

The multichannel anisotropic filter can be formulated as $\frac{\partial}{\partial t}\overrightarrow{I}(\overline{x},t) = \operatorname{div}\left(c_c(\overline{x},t) \times \nabla \overrightarrow{I}(\overline{x},t)\right)$, where $c_c(\overline{x},t) = f(\left\|\nabla \overrightarrow{I}(\overline{x},t)\right\|)$, replacing the $|\nabla I(\overline{x},t)|$ term in the diffusion function by the Euclidean Norm $\left\|\nabla \overrightarrow{I}(\overline{x},t)\right\|$.

5 Results

In this section we will examine the results obtained running the proposed algorithm on a range of images.

The noise tolerance of our segmentation method can be demonstrated by looking at its effect on an artificial 2D image, given by

$$g(x,y) = \begin{cases} -10 & x \in [-50,-10) \\ x & x \in [-10,10) \\ 10 & x \in [10,50) \end{cases}$$

for $y \in [0, 100)$. In this image, we have regions with a difference in graylevels of 20, separated by a broad transition region (Figure 1a). The segmentation result on this image is a vertical line at $x = 0$ (Figure 1b), as we have expected. To test the stability of URG with the presence of noise, we have added Gaussian noise with zero mean and standard deviation of $\sigma = 2.5$, $\sigma = 5.0$ and $\sigma = 10.0$ (Figures 1c, 1e and 1g respectively). The segmentation results are shown in Figures 1d, 1f and 1h. As can be seen, the approximate location of the boundary has been detected correctly even in the presence of high level of noise.

Figure 2a consists of uniform radial gradient on top of linear gradients. The segmented image is shown in Figure 2b. It can be seen that the radial and linear divisions are relatively uniform, and that the decision of region boundaries are consistent across disjoint areas of the image. In essence, the URG procedure is not free of scan order dependencies, but it is demonstrated that the impact is minimal.

URG has also been applied to a range of medical images with remarkable results.

Figure 3 shows the vessel segmentation of a low contrast X-ray angiogram image. In Figure 3a, the contrast varies between different parts of the image, and the intensity of vessels are not uniform. In addition, a relatively large noise gradient is present, making the segmentation of fine vessels a challenging task. Figure 3b shows the segmented vessels. Even though several fine vessels were missed, most of the major vessels have been detected.

Figure 4 shows ventricular space segmented from a low resolution ultrasound heart image. The image in Figure 4a contains speckle artifacts commonly found in ultrasound images. The large middle region in the segmented image represents the ventricular space, and its outline is displayed in Figure 4b, overlaying on top of the original image.

Figure 5 shows the segmentation of cells in a low contrast microscopic image. The image in Figure 5a has been corrupted by high frequency noise, so it is essential for the segmentation algorithm to be noise tolerant in order to achieve satisfactory results. Figure 5b displays the detected cells, which are useful for cell counting purposes.

(a)

(b)

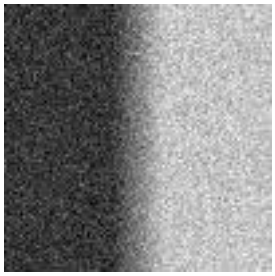

(c)

(d)

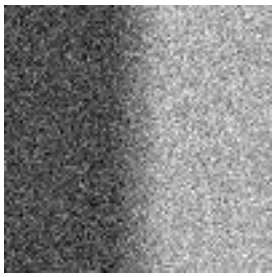

(e)

(f)

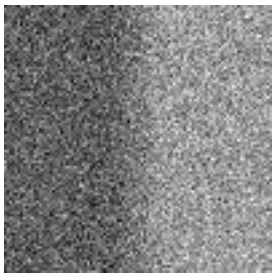

(g)

(h)

Figure 1: Segmentation of artificial images with additive noise

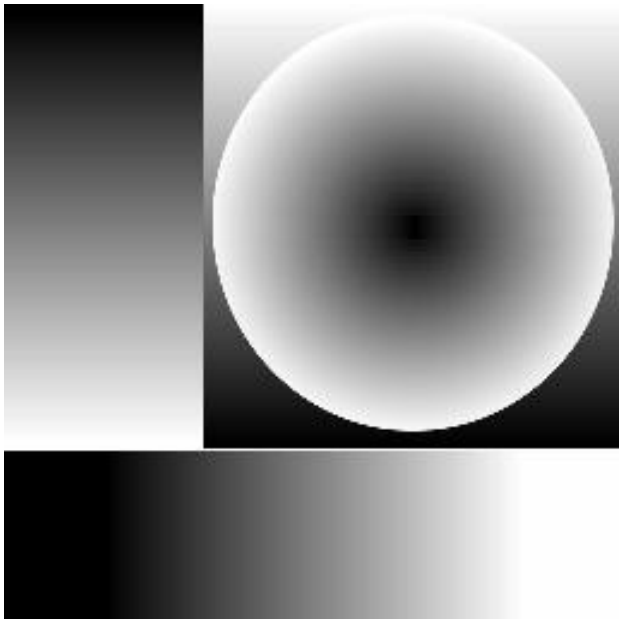

(a)

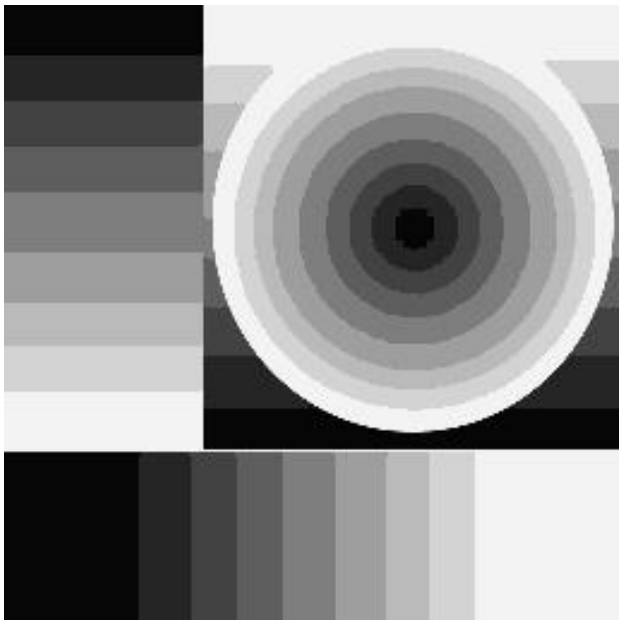

(b)

Figure 2: Segmentation of gradient image

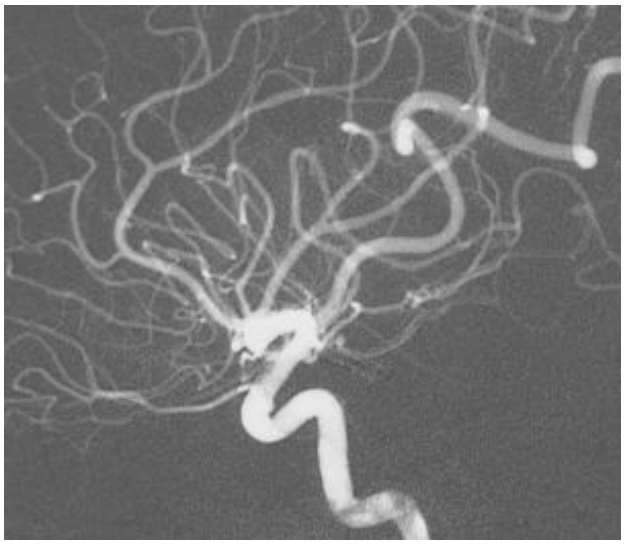

(a)

(b)

Figure 3: Segmentation of X-ray angiogram

(a)

(b)

Figure 4: Segmentation of ultrasound heart image

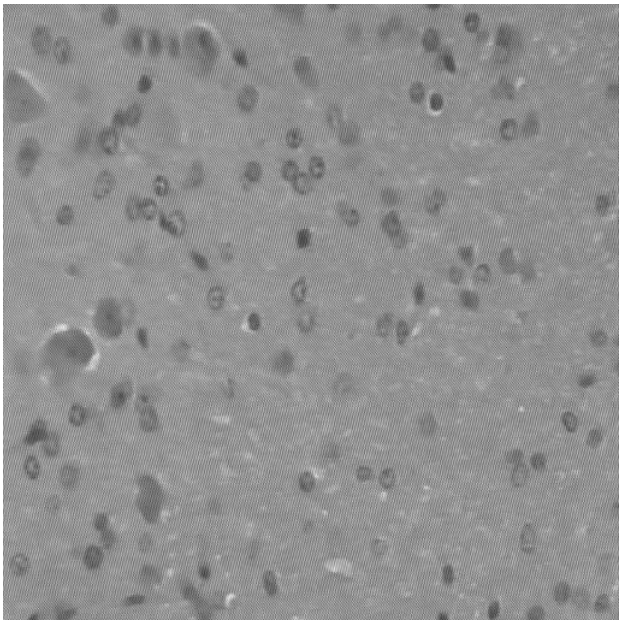

(a)

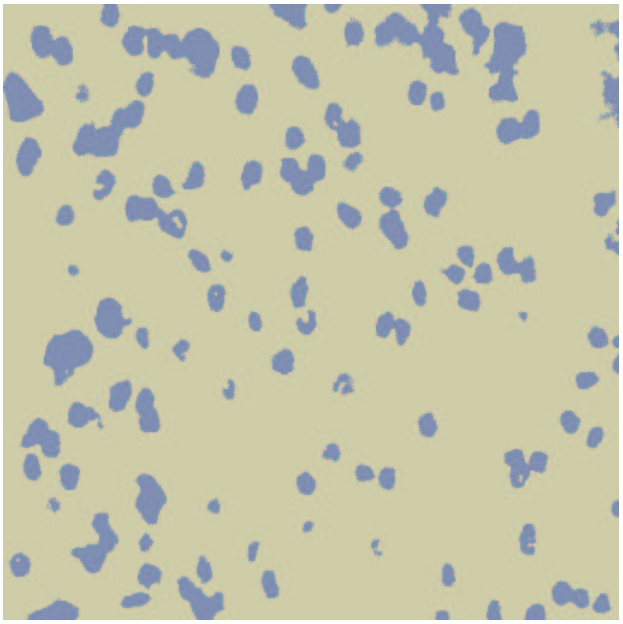

(b)

Figure 5: Segmentation of microscopic cell image

References

[1] R. Adams and L. Bischof. Seeded region growing. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 16(6):641–647, 1994.

[2] R. Beare and H. Talbot. Exact seeded region growing for image segmentation. In *Proc. DICTA, 5th Bienn. Conf. Aust. Patt. Rec. Soc*, pages 132–137, December 1999.

[3] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 10(2):167–192, 1988.

[4] L. S. Davis. A survey of edge detection techniques. *Comput. Graphics Image Process.*, 4:248–270, 1975.

[5] G. Gerig, O. Kbler, R. Kikinis, and F. A. Jolesz. Nonlinear anisotropic filtering of mri data. *IEEE Transactions on Medical Imaging*, 11(2):221–232, June 1992.

[6] R. M. Haralick and L. G. Shapiro. Image segmentation techniques. *Comput. Vision Graphics Image Process.*, 29:100–132, 1985.

[7] S. L. Horowitz and T. Pavlidis. Picture segmentation by a directed split-and-merge procedure. In *Proc. 2nd. Int. Joint Conf. on Pattern Recognition*, pages 424–433, 1974.

[8] M. Kass, A. Witkin, and D. Terzonpoulos. Snakes: Active contour models. In *Proc. Int. Conf. on Computer Vision*. London, 1987.

[9] F. Meyer and S. Beucher. Morphological segmentation. *Journal of Visual Comm. and Image Representation*, 1:21–46, 1990.

[10] N. Nordstrm. Biased anisotropic diffusion - a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8(4):318–327, 1990.

[11] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 12:629–639, 1990.

[12] P. K. Sahoo, S. Soltani, and A. K. C. Wong. A survey of thresholding techniques. *Comput. Vision Graphics Image Process.*, 41:230–260, 1988.

[13] Y. Wang, J. S. Jin, and J. Hiller. An adaptive nonlinear diffusion algorithm for image filtering. In *Proc. SPIE: Real-time Imaging*, volume 3028, pages 26–37, 1997.