



Multi-class 3D region growing algorithm

Jakub Smółka*

*Department of Computer Science, Faculty of Electrical Engineering and Computer Science,
Lublin University of Technology, Nadbystrzycka 36b, 20-618 Lublin, Poland*

Abstract

This paper describes generalization of multi-class region growing algorithm allowing for segmentation of 3D images (series of slices). The multi-class region growing algorithm was proposed in [1]. Additionally, a new method for finding the start region was presented. As its 2D version the new algorithm does not need initial parameters since it features segmentation quality assessment. A series of segmentations is performed on a dataset, each segmentation quality is assessed and the best one is picked. Additionally, the number of classes in the image is determined automatically. The multi-class 3D region growing algorithm is tested on CT and MRI scans. Different types of MRI scans are used. The scans come from multiple sources. The results are shown in the form of 3D reconstruction accompanied by a selected 2D slice. In addition, such selected algorithm performance issues are discussed as effective algorithm implementation; fast dilation implementation is also mentioned. The paper explains all concepts and operations used by the algorithm. It includes numerous figures and algorithm pseudo-code descriptions.

1. Introduction

With developments in CT and MR imaging, the acquisition of high quality images is possible. However, in order for these images to be an effective support in diagnosis or quantitative analysis, they must be processed to extract relevant parts of the anatomy [2,3]. Segmentation is a very important step, which determines the accuracy of subsequent processes. Segmentation algorithms can be divided into two classes: (1) algorithms that work with 2D images (2) algorithms that work with 3D images. The second class is usually computationally more expensive. However, one of their advantages is that they treat series of slices as a whole; they do not segment each image separately. Moreover, the rapidly increasing computing power of modern computers, along with program optimization, allows for the effective use of 3D segmentation algorithms.

* E-mail address: jakub.smolka@pollub.pl

2. Definitions

A *structuring element* can be thought of as a 3D filter mask $S(x,y,z)$ in which coefficients take binary values [4]. These values specify desired shape the origin neighborhood. In the described algorithm, a cubical neighborhood (Figure 1a) consisting of 27 voxels is used (all pixels in the $3 \times 3 \times 3$ structuring element are ON).

Dilation is used by the described region growing algorithm for enlarging the region by adding ON pixels at the region boundaries. Figures 1b and 1c show an example of 3D region before and after dilation. When a structuring element is over a region, an OFF pixel at the origin is set to ON if any of the structuring elements overlaps ON pixels of the image [4-6].




Fig. 1. Neighborhood shape used in the 3D multi-class region growing (a), example 3D region (b), example 3D region after dilation (c)

Homogeneous region is a region whose standard deviation σ is less or equal to a given homogeneity threshold σ_{max} [7,2].

Histogram contraction is used during the region growing process for restoring homogeneity of the grown region. This is achieved by performing a thresholding operation that removes voxels corresponding to the k extreme non-void histogram bins. *K-left-contraction* or *k-right-contraction* is used by the region growing algorithm depending upon which objects, bright or dark, are segmented.




Fig. 2. Examples of 4-left contraction (a) and 4-right contraction (b)

3. 3D Region Growing Algorithm

The region growing algorithm, which has been implemented, is a modified and simplified version of the algorithm proposed by Revol and Jourlin [7]. It is capable of segmenting non-connected regions. The voxels, which belonged to a homogenous region at a certain step, can be removed later. This property allows it to segment a complex region, not requiring a seed point in all of its parts. The modification, mentioned above, allows the algorithm to work only on a selected part of the image. This is achieved by performing a logical AND operation on two binary images. These images describe a dilated region and the selected part of the image.

At each step a 3D homogeneous region R_n is dilated, yielding a region R_{n+1} . If R_{n+1} is homogenous, the algorithms proceed to the next iteration. Otherwise it reduces R_{n+1} histogram dynamic range by performing 1-left-contractions or 1-right-contractions until R_{n+1} is homogeneous and the algorithm can proceed to the next step. When the object is brighter than the background the left-contraction is used; otherwise, the right-contraction is used.

One of the advantages of this algorithm is that 2D implementation can be easily converted to 3D. The only thing that has to be done is reimplementing the basic operations used by the algorithm such as dilation, k-contraction and logical operations (for regions) in 3D.

The following presents a pseudo-code description of the region growing process used.

```

n ← 1;
R0 ← 0; //one before last = empty
R1 ← Ms; //last = seed points
while (Rn ≠ Rn-1) // continue if regions are different
    Rn+1 ← Dilation3D(Rn);
    Rn+1 ← Rn+1 and Mb; //limiting the region
    if (σRn+1 ≥ σmax) //is Rn+1 not homogenous
        do //contracting the histogram
            if (Object=bright)
                Rn+1 ← LeftContraction3D(1, Rn+1);
            else Rn+1 ← RightContraction3D(1, Rn+1);
        while (σRn+1 > σmax) // is Rn+1 homogeneous?
    end;
    n ← n+1; //next iteration
end;
RG ← Rn; //result

```

A straightforward implementation of the algorithm shown above is computationally ineffective. During one iteration, histogram contractions are

usually performed a few hundred times. After each histogram contraction the standard deviation has to be calculated. This can be lengthy for the volumes containing a few million voxels. A much better approach is to build a histogram of dilated region, to find the optimal threshold for histogram contraction by removing only histogram bins (not corresponding pixels) and to calculate the standard deviation based only on the modified histogram. After the threshold is found, only one actual thresholding operation is performed, giving a homogeneous region.

The region growing algorithm can be optimized even further since it uses another computationally expensive operation - 3D dilation. Dilating a $256 \times 256 \times 100$ region can take up to a few seconds with a simple 3D dilation implementation. Using highly optimized implementation for binary images can reduce these times by a factor of 20 [8]. This is important because there might be even a few hundred dilations during one region growing process.

4. Automating segmentation

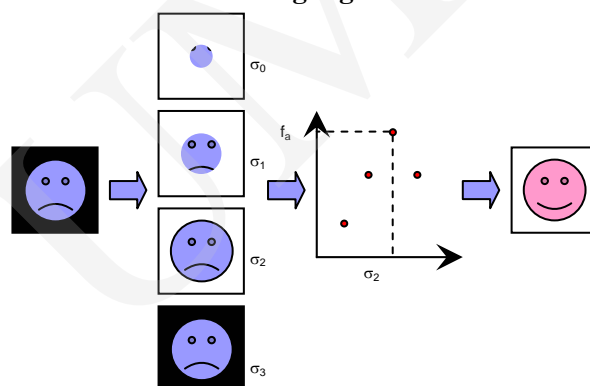


Fig. 3. Principle of an automated region growing

The 3D region growing algorithm described above grows a homogeneous region from a given start region. In order to do that it needs a homogeneity threshold as one of its parameters. It is difficult to determine this threshold *a priori*. A solution to this problem was introduced by Revol, Peyrin, Carrillon and Odet in [2]. An automated region growing process consists of image segmentation, segmentation quality assessment and best segmentation selection (figure 3). First an image is segmented for each σ_{max} value within a given range ($\sigma_0 - \sigma_3$). Each time the segmentation starts with the same start region. After each region growing process, the segmentation quality is assessed and the σ_{max} value is increased by $\Delta\sigma$ until it reaches σ_3 . Finally the optimal value of the standard deviation is determined and, therefore, the optimal segmentation

obtained. In this approach the maximum value of the assessment function corresponds to the best segmentation in the sequence.

Please note that the concept of quality assessment remains unchanged regardless of the dimensionality of segmented image. The only thing that had to be reimplemented for 3D version of the algorithm is the assessment functions.

5. Segmentation quality assessment

Quality assessment is crucial to automating the region growing process. Not only does it eliminate the need for a supervisor, but it also guarantees repeating results for the same input data.

Two basic approaches to quality assessment are possible: by boundary or by region [2]. In the first case the best segmentation corresponds to the regions whose boundaries present the strongest edge. In the second case the best segmentation corresponds to the most homogenous region in the original image. In the boundary approach many different assessment functions are possible, including the functions based on a logarithmic image processing model [9], standard deviation or transition levels. In the region approach, types of assessment functions include the functions based on entropy, distance from the gray level function (image), and inter-cluster variance. Some of these functions were derived from automatic thresholding methods such as Otsu's (inter-cluster variance) method or the entropy maximization method.

All of the above mentioned methods were tested on 3D images. Only the logarithmic image processing contrast and maximum entropy methods gave good results. However, since the results for the maximum entropy method were slightly better, only this method is introduced in this paper.

5.1. Maximum entropy

Entropy is commonly used for optimal image thresholding [4]. This method consists in maximizing the sum of the entropies of the two classes into which the automated region growing divides the image. If \bar{R} is the complement (background) of R (object) and p_{iR} , $p_{i\bar{R}}$ are the probabilities of individual gray levels, then the sum of the entropies is defined as follows:

$$S(\sigma_{\max}) = -\sum_{i=0}^{M-1} p_{iR} \ln p_{iR} - \sum_{i=0}^{M-1} p_{i\bar{R}} \ln p_{i\bar{R}}$$

where M is the number of available gray levels (for example $M = 4096$ if each voxel is described using 12 bits).

Note that voxels having the same value may be both in the object and the background.




Fig. 4. Histograms of region and its background

6. Multi-class region growing

Typically in medical images, there are more than two structures; therefore, the segmentation algorithm should be able to segment more than two classes. However, the automated region growing introduced by Revol, Peyrin, Carrillon and Odet in [2] can segment only two classes – object and background. This limits its application to the images with bimodal histogram.

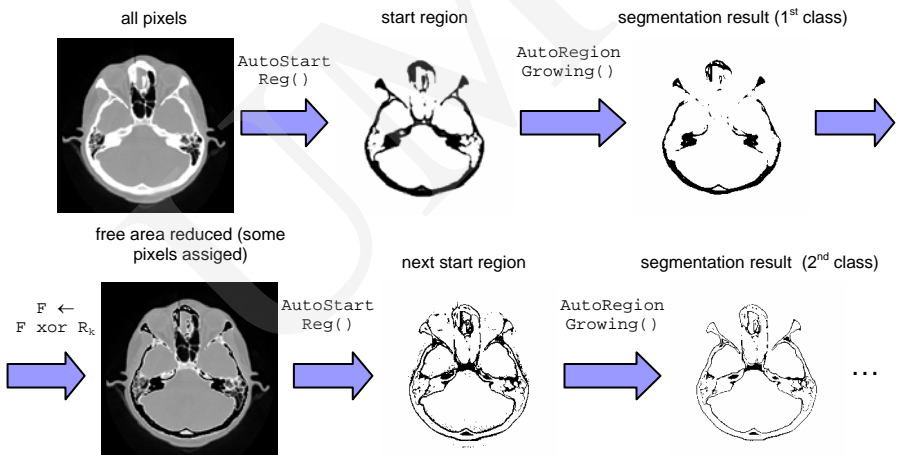


Fig. 5. Concept of multi-class region growing

In order to circumvent this limitation, the 3D multi-class region growing is introduced. This new algorithm is, in fact, a sequence of automated region growing processes. At the beginning of multi-class segmentation, a start region is automatically determined and its standard deviation is calculated. The calculated value is then used to determine the range which will be searched for the optimal homogeneity threshold. Selecting a range from 0.9 to 1.1 of the start region standard deviation usually gives the best results for 3D images. The automated region growing is started. As a result the brightest class in the picture is obtained. Then the voxels that belong to the first class are marked as assigned so that they cannot be assigned to any other class. A new start region is obtained

and the whole sequence is repeated. Multi-class segmentation stops when no further class can be segmented.

The following presents a pseudo code description of the multi-class region growing.

```

 $n_0 \leftarrow 0;$ 
 $F \leftarrow M_b;$  //all voxels unassigned
 $n_1 \leftarrow \text{OnPointCount3D}(F);$  //number of unassigned voxels
 $k \leftarrow 1;$ 
while ( $n_k \neq n_{k-1}$ ) //has a new class been segmented?
   $M_t \leftarrow \text{AutoStartReg3D}(\text{MinVolume}, F);$  //new start region
  if ( $\sigma_{\text{const}} = \text{false}$ )  $\sigma_r \leftarrow \sigma(M_t);$  //calculate  $\sigma$  of new
    //start region
     $R_k \leftarrow \text{AutoRegionGrowing3D}(0, 9 \cdot \sigma_r, 1, 1 \cdot \sigma_r, M_t, F);$ 
   $F \leftarrow F \text{ xor } R_k;$  //mark voxels as assigned
   $k \leftarrow k+1;$ 
 $n_{k-1} \leftarrow n_k;$  //save the number of previously
  //unassigned voxels
 $n_k \leftarrow \text{OnPointCount3D}(F);$  //new number of unassigned
  //voxels
end;
MCRG  $\leftarrow R_1, R_2, \dots, R_{k-1};$  //segmentation result

```

Note that the concept of multi-class region growing is general enough that it can be used with any other two-class automatic segmentation method with only minor modifications in the algorithm initialization part. Additionally, enabling the algorithm to work with 3D images required only reimplementing of the operations it uses, such as finding the start region.

7. Obtaining the start region

Apart from the homogeneity threshold, each region growing process requires a start region. Automated region growing requires only one start region since each region growing process starts with the same set of seed points. However, multi-class region growing needs a new start region for each class it tries to segment, so a way of obtaining it automatically is necessary. Experiments have shown that an automated region growing result remains, in many cases, unaffected by changes in the start region. Therefore, a simple method is used in the multi-class region growing. The method used selects the brightest region (consisting of unassigned voxels), whose volume is just above the minimum set, arbitrarily (Figure 6). 1/10 of the total number of voxels in the picture usually gives the best results.




Fig. 6. Obtaining the start region

8. Results

The 3D multi-class region growing algorithm has been tested on CT and MRI scans. The following presents the selected results. All input images were smoothed (slice by slice) prior to segmentation with a 3×3 Gaussian filter. Each automatic region growing process was initialized with a start region containing 1/10 of the total number of voxels in the volume. The standard deviation range which was searched for the optimal homogeneity threshold was limited to [0.9,1.1] of the start region standard deviation. In all presented segmentations the maximum entropy method was used for quality assessment.

Figure 7 presents segmentation of a CT scan. Multi-class region growing has distinguished 4 classes. Two of them belong to the background, and the other are soft tissue and bone.




Fig. 7. Segmentation of a CT scan series (CT data set included with VTK 4.2 library was used)

Figure 8 compares automatic segmentation of a series of MRI scans with it is manual segmentation performed by an expert. Only two classes could be compared, the others were manually removed from the automatic multi-class segmentation.




Fig. 8. Comparison of 3D multi-class region growing segmentation with the expert manual segmentation (MR brain data set 788_6_max and its manual segmentation was provided by the Center for Morphometric Analysis at Massachusetts General Hospital and is available at <http://neuro-www.mgh.harvard.edu/cma/ibsr.>)

Figure 9 depicts 3D reconstruction of 3 classes that were distinguished by the 3D multi-class region growing algorithm.




Fig. 9. 3D reconstruction of a brain (3 classes) (m_vm10xx data set from the Visible Human Project was used)

9. Conclusions

The concept of multi-class region growing is general enough to be implemented both in 3D and in 2D. Changing the dimensionality of the algorithm requires only reimplementing of basic operations such as dilation or histogram contraction. Experiments have shown that 3D multi-class region growing is able to deliver good quality segmentation for many data sets using default parameters. The quality is comparable to that of segmentations performed by an expert. That is why, in most cases, the algorithm does not require the user input. Segmentation of $256 \times 256 \times 64$ volume takes only 3-4 minutes (on Athlon XP 2500+ computer) whereas segmentation of 256×256 image takes 3-5 seconds. This means that relation between the computation time and the amount of segmented data is approximately linear. Thanks to the introduced optimizations the change in dimensionality had no impact on computation time. The results obtained with 3D version of the multi-class region growing are similar to those obtained with its 2D version. However, the 3D segmentation allows for avoiding main pitfalls of segmenting 2D slices separately. The segmented classes are coherent and marked using the same labels. Experiments have shown that not all assessment methods are suitable for 3D multi-class region growing. The best results were obtained using the maximum entropy method. With it the algorithm was able to grow coherent regions, and properly detect number of classes.

References

- [1] Smolka J., *Multi-class Region Growing Algorithm*, Annales Informatica UMCS, Lublin, 1 (2003) 165.
- [2] Revol-Muller Ch., Peyrin F., Carrillon Y., Odet Ch., *Automated 3D region growing algorithm based on an assessment function*, Pattern Recognition Letters, 23 (2002) 137.
- [3] Suri J, Setarehdan S, Singh S, *Advanced Algorithmic Approaches to Medical Image Segmentation*, Springer-Verlag London Limited, (2002).
- [4] Seul M., O’Gorman L., Sammon M.J., *Practical Algorithms for Image Analysis: Description, Examples, and Code*, Cambridge University Press, (2000).
- [5] Myler H.R., Weeks A.R., *The Pocket Handbook of Image Processing Algorithms in C*, Prentice Hall PTR.
- [6] Gonzalez R.C., Woods R.E., *Digital Image Processing*, Addison-Wesley Publishing Company, (1993).
- [7] Revol Ch., Jourlin M., *A new minimum variance region growing algorithm for image segmentation*, Pattern Recognition Letters, 18 (1997) 249.
- [8] Smolka J., *Szybka realizacja podstawowych operacji morfologicznych*, Obliczenia naukowe – Wybrane problemy, Lublin, (2003) 93, in Polish.
- [9] Pinoli J., *A general comparative study of the multiplicative homomorphic, log-ratio, and logarithmic image processing approaches*, Signal Processing, 58 (1997) 11.