

# Dynamic Region Growing

Andreas Siebert

Department of Computer Science  
University of British Columbia

## Abstract

A novel segmentation algorithm based on the region growing paradigm is presented. Unlike previous segmentation methods, this novel scheme requires neither hand-tuning of parameters nor knowledge about the scene. Instead, the parameter which controls the segmentation is dynamically derived from the data for each region, based on a local quality measure of the region's contour. The algorithm assigns a stability value to each extracted region that reflects the robustness of that region. Results are shown for some gray level aerial images.

## 1 Introduction

We are developing a segmentation algorithm as a first step in an image indexing and retrieval system. Some important design criteria for such a segmentation algorithm are as follows:

- The algorithm must deliver homogeneous *regions* from which various region based features (color, shape, texture, etc.) can be computed. Edges that don't form closed contours, as for example produced by the Canny edge detector [2], don't lead directly and naturally to regions (contour *gap* problem). Edge detectors that by their design produce closed contours, as for example the zerocrossings of the Laplacian of Gaussian, usually don't deliver the desired regions, especially in noisy images.
- Hand-tuning of parameters, e.g. of thresholds, is not acceptable. The system should work fully automatically.
- A measure to quantify the stability of each region must be provided.
- Because of the wide variety of images in a typical image database, no assumptions can be made about the image in terms of lighting conditions, objects to be expected, scale, number or shape of regions etc. That is, we cannot assume any

domain knowledge or the existence of object or scene models.

The latter point signifies a departure from the approach most image processing systems take these days. Certainly, better results can be expected if we can integrate knowledge about the expected objects, the lighting conditions, the scene geometry, etc., but in the application we have in mind, we are much more likely to encounter images where we lack prior knowledge but still want to perform some useful actions. Therefore, it is important that the algorithm can measure the quality of the segmentation not in terms of model selection and parameter fitting but in terms of the image data itself. We propose a metric that measures the strength of a contour relative to its neighboring pixels which are not part of the contour. Since the algorithm dynamically adjusts the merging threshold for each region in order to maximize the contour strength, we call it *Dynamic Region Growing* (DRG). DRG addresses not only the automatic setting of thresholds but also the unwanted region chaining problem that has plagued many region growing algorithms [5].

Some classic region growing techniques have been described and classified in [5]. Their greatest drawback has been their dependence on a good choice of the parameters involved. A recent review [10] includes references to segmentation methods based on fuzzy logic, neural networks, and color. A seeded region growing method has been proposed in [1]. The authors claim that their algorithm does not require parameter tuning; however, it critically relies on the seed points being given as input. Active contour models ("snakes") [7] fit contours with energy-minimizing splines. These are attracted to nearby edges, but they have to be placed somewhere near the desired contour by external forces, which makes them useful in an interactive environment rather than in a fully automatic system. [15] aims at a unification of region growing with active contour models and Bayesian techniques that allow the integration of global constraints. Problems of global approaches with respect to an application in image indexing and retrieval include the

exhaustiveness (each pixel must be labeled) and exclusiveness (no overlapping regions) constraints of the segmentation, the adequate choice of a family of probability distributions (a Gaussian distribution too often is the ad hoc choice), and the slow convergence of the optimization.

Many schemes have been designed to integrate region growing algorithms with edge detection in order to get the best of both worlds. [3] uses a maximum likelihood estimator to merge different edge maps, followed by region growing to satisfy some constraints on the desired regions. [4] proposes an involved procedure to determine seed points (called *germs*); the region growing itself is controlled by the magnitude of the gradient at each pixel.

Ohlander's recursive region splitting method [9] was an early attempt to dynamically compute the segmentation parameters from the image data. The algorithm starts with the whole image being one region which is then recursively split. The decision whether and how to split a region is based on the shapes of the histograms of the respective feature values. Note that histograms don't keep spatial information.

## 2 Algorithm

In the absence of object or scene models, it is critical to have a criterion that enables us to evaluate the quality of a segmentation. We choose this criterion to be the *contour strength* where we define the contour strength  $cs(R)$  of a region  $R$  to be the sum of the absolute differences between each pixel on the contour of a region and the pixels in the 4-neighborhood of these contour points that are not part of the region under consideration, i.e.

$$cs(R) = \frac{1}{n} \sum_{p_i \in C_R} |p_i - q_i| \quad (1)$$

where  $C_R$  is the set of pixels on the contour of  $R$ ,  $q_i \notin R$  is in the 4-neighborhood of  $p_i$ , and  $n$  is the number of pixels  $p_i \in C_R$ . The general idea is that regions are bound by strong contours, so we grow regions such as to maximize  $cs_R$ . Note that contour following and taking differences of integer numbers are computationally inexpensive operations. Slightly improved results at higher computational costs can be expected if the contour strength is based on the intensity *gradient* at each contour pixel rather than on the intensity difference.

Traditional region growing algorithms base their decision whether to merge a region with a pixel on a homogeneity criterion that usually takes into consideration the mean and standard deviation of the regions [1],[5]. DRG is similar in that it grows regions

based on the difference between the mean of a region and the intensity of the pixel to be merged. However, DRG maximizes the contour strength which is a purely local criterion that allows pixels to be merged with a region if this merging results in a stronger contour, even if those pixels don't meet such homogeneity criteria. We contend that our criterion is more useful in the presence of noise since it does not depend on any a priori homogeneity thresholds or statistics.

To be more specific, let  $R_\epsilon$  be a region to be grown, and  $\mu_\epsilon$  the mean intensity of that region. The region growth is controlled by a parameter  $\epsilon$ : A pixel  $p$  with intensity  $I_p$  that is spatially adjacent to the region  $R_\epsilon$  is merged with  $R_\epsilon$  if  $|\mu_\epsilon - I_p| \leq \epsilon$ .  $\mu_\epsilon$  is updated immediately.

For each region, DRG starts with a seed that consists of one or more pixels.  $\epsilon$  is set to a low value  $\epsilon_{low}$ . Then DRG loops through increasing values of  $\epsilon$  until a termination criterion is met. For each  $\epsilon$ , we compute the corresponding contour strength  $cs_{R_\epsilon}$  and keep the maximum strength. In other words, for each seed DRG creates a whole family of segmentations and then picks the best one. In order to define an adequate termination criterion, we must address the problem of *overspill*.

Region growing algorithms generally are subject to overspill, referred to as *region chaining* in [5]. Overspill occurs when two regions that should be separate are merged. An example is given in figure 1. Overspills are frequently caused by noise and therefore happen quite often in natural images.




Figure 1: Left: before overspill; Right: after overspill

Generally, it is undesirable to continue region growing beyond an overspill. Overspills are not apparent from looking at the contour strength alone. However, detection of an overspill is fairly straightforward and robust. An overspill is usually characterized by the sudden emergence of three properties: (i) a large increase in the region size, (ii) a decrease of the circu-

larity of the region, and (iii) a change in the average region intensity. We define circularity to be the ratio of region size to the square of the contour length. Typically, an overspill creates a small “bridge” between the two regions chained together which in turn causes a large increase in the contour length and therefore a decrease in circularity. These three properties are combined in a fuzzy function to decide whether an overspill has occurred.

We can now return to the problem of determining the range of the parameter  $\epsilon$ . DRG sets  $\epsilon_{low}$  to a small integer value. The more noise in the image, the higher  $\epsilon_{low}$  should be. The exact value of  $\epsilon_{low}$  is not critical. The conservative strategy is to set  $\epsilon_{low}$  low enough such that no good segmentation is to be expected for  $\epsilon < \epsilon_{low}$ . As for the upper bound of  $\epsilon$ , DRG stops either when an overspill has been detected or when an upper bound  $\epsilon_{high}$  has been reached. Again, the exact value of  $\epsilon_{high}$  is not critical. The conservative strategy is to set  $\epsilon_{high}$  high enough such that no good segmentation is to be expected beyond it. We implemented DRG such that for each region at least three different values of  $\epsilon$ , i.e.  $\epsilon_{low}$  to  $\epsilon_{low} + 2$ , are tried.

A high level description of algorithm DRG is given in figure 2. It remains to elaborate on the strategies for the selection of seed points and for dealing with overlapping regions.

```

while seeds remain
  find next seed  $R_{i,0}$ 
   $\epsilon = \epsilon_{low}$ 
  while  $\epsilon \leq \epsilon_{high}$  and no overspill
    grow region  $R_{i,\epsilon}$ 
    compute contour strength  $cs_{R_{i,\epsilon}}$ 
     $\epsilon = \epsilon + 1$ 
  end while
end while
 $R_i = R_{i,\epsilon} \mid (cs_{R_{i,\epsilon}} = \max_{\epsilon}(cs_{R_{i,\epsilon}}))$ 
if  $R_i$  does not coincide with  $R_j, 1 \leq j < i$ 
then keep  $R_i$ 

```

Figure 2: Dynamic Region Growing DRG

Several strategies for choosing the seed points are conceivable. The lack of scene knowledge implies that we cannot assume a fixed number or certain initial positions of the seeds. Instead, they must be derived from the data. The strategy we chose for DRG is this: find the next *anchor* that is not labeled as being part of a region extracted so far. By anchor, we mean a homogeneous micro-region where all pixel intensities lie within  $\epsilon_{low}$  of the intensity of the upper left pixel

of that micro-region. This way, DRG does not try to grow regions from noisy pixels or areas that are obviously not homogeneous. The adequate anchor size depends on the image. If the image is smooth, a larger anchor size could be chosen. Under the no-knowledge assumption, the conservative strategy used by DRG is to chose a small anchor size. Note that the computational workload is linear in the number of seeds. Therefore, reducing the number of seeds is an effective way to reduce the runtime of DRG, at the cost of possibly missing some regions with high contour strength.

The seed picking strategy adopted assures that there are no identical regions – they must differ at least in the seed pixels of the second region. However, they might overlap a lot. Unlike most other region growing schemes, DRG does not prevent a region from growing into an area already labeled by another region. The reasons behind it is that (i) it is not necessarily true that the previous labeling is better than the new one, and (ii) for many applications, having mutually exclusive regions is not a requirement. In fact, mutual exclusion might turn out to be too hard a constraint given the assumptions above. It has been argued that regions might naturally overlap [11], e.g. “sky” and “tree”. Furthermore, segmentations based on other criteria than homogeneity or contour strength, e.g. texture based segmentation, will lead to regions that also overlap with the regions extracted by DRG. Let us define two regions to coincide if they have more than 90% of their pixels in common. The simple approach adopted in this paper is to eliminate the region with weaker contour strength if two regions coincide, otherwise to keep both. In the present implementation, DRG also eliminates all regions below a certain size, typically 100 pixels.

As  $\epsilon$  increases, the size of the corresponding region  $R_{i,\epsilon}$  either also increases or remains constant. A straightforward measure to evaluate the *stability* of a region is to determine the range of  $\epsilon$  such that  $R_{\epsilon}$  is similar in size to the region  $R_{\epsilon_{max}}$ . Therefore, we define the stability  $stab(R)$  of region  $R$  to be

$$stab(R) = \epsilon_{top} - \epsilon_{bottom} \quad (2)$$

where

$$\epsilon_{top} = \max(\epsilon) \mid \frac{size(R_{\epsilon}) - size(R_{\epsilon_{max}})}{size(R_{\epsilon_{max}})} \leq \theta$$

$$\epsilon_{bottom} = \min(\epsilon) \mid \frac{size(R_{\epsilon_{max}}) - size(R_{\epsilon})}{size(R_{\epsilon_{max}})} \leq \theta$$

Again, the choice of  $\theta$  is not critical, and  $\theta$  remains constant for all images. Here, we set  $\theta = 0.2$ , i.e. we

look at the range of  $\epsilon$  for which the region sizes haven't changed by more than 20%.

Note that a high stability implies a high contour strength, but the converse is not generally true. Recall that the contour strength sums up all individual intensity differences along the contour line. Even if the overall sum is large, there might be some "weak links" in the "contour chain", resulting in an overspill and therefore in a low stability. By contrast, a high stability reveals that there is no such weak link along the contour line.

Our notion of stability is not unlike Witkin's definition of stability as the range of a scaling parameter  $\sigma$  over which no visual events occur [14].

### 3 Properties

The four most important properties of DRG follow directly from the design criteria above:

- DRG returns spatially coherent regions whose contour lines form closed curves.
- No hand-tuning of parameters is required. The most critical parameter,  $\epsilon$ , is automatically varied over a certain range. DRG employs a contour strength criterion to select the best value for  $\epsilon$ . This best  $\epsilon$  is determined for each region separately.
- DRG provides a stability measure for each region extracted.
- DRG does not use any domain knowledge.

Other properties of DRG are as follows:

- The regions extracted can overlap, i.e. the segmentation is not exclusive. Also, the union of all regions does not necessarily cover the whole image, i.e. the segmentation is not exhaustive.
- The regions extracted might contain "holes". Closing such holes, if required by an application, is left to a postprocessing stage. Note that the contour strength of a region is exclusively based on the *outer* contour line, not on any *inner* contour lines caused by holes.
- The region contours tend to be rugged. DRG does not attempt to smooth boundaries. While smooth boundaries are a desirable feature in many domains, especially if it is known that many man-made structures are likely to occur in an image, such a smoothing cannot be justified in the absence of such domain knowledge. However, if

smooth boundaries are required, then the approach given in [13] which combines regions with active contour models can be used to obtain them.

- DRG uses a purely local criterion, based only on the contour pixels and their 4-neighborhoods, to evaluate a region. No global constraints are considered.
- DRG provides a stability measure for each region extracted.
- DRG is not guaranteed to find the optimal solution. Basically, there are three reasons for this. First, DRG grows regions by incrementing  $\epsilon$  in steps of one. This quantization might cause the optimal solution to be missed. Second, the strategy chosen for selecting seed points might miss some seeds necessary to obtain the optimal solution. Third, DRG is not completely independent of the scanning direction.
- DRG offers an intuitive way to measure a region's stability.
- A generalization to multispectral images is straightforward. Only the distance function for computing the distance between a region's mean vector and a pixel's intensity vector has to be modified. Typically, a Euclidean norm is used. An alternative approach is to run DRG separately on each band and then to combine the results.
- DRG shares some less desirable properties with other region growing algorithms. Most notable are the dependence of the results on the order in which pixels are added to a region, and the difficulty of a precise mathematical analysis because of the inherent nonlinearity of region growing. On the other hand, DRG successfully addresses the problem of region chaining.

### 4 Experimental Results

Figure 3 shows a 160×140 pixels aerial image. It contains various fields and roads. Figure 4 shows the result of applying DRG to the original image, while figure 5 shows only those regions of the previous image that show a high degree of stability ( $stab(R) \geq 4$ ).

In the example given in figure 4, an arbitrary lookup table has been used to render the 27 extracted regions. However, a compact visualization of the results is difficult because regions can and do overlap. Here, in the printout, regions with a larger label simply overwrite regions with a smaller label at pixels

where regions overlap. Unlabeled pixels are shown in black.

Some of the properties of DRG mentioned above, like rugged region boundaries and holes in the regions, are clearly visible in the segmentation shown in figure 4. At first sight, it may seem desirable to have smooth contour lines of the fields along the roads, but the intensity values don't justify such a segmentation. Also, it seems surprising that the large road across the upper left part of the image is not labeled at all. However, a closer look reveals that the intensity of the road varies dramatically; it is not a very homogeneous region.

Fig. 6 shows the result obtained by running DRG on a rotated version (90 degrees) of the original image. Some of the stable regions are almost identical to the ones extracted from the original image, but others are not, demonstrating the undesirable dependence of region growing algorithms on the scanning direction.

We looked also at the behavior of DRG with respect to noise. Fig. 7 shows the stable regions extracted from the original image with gaussian noise added (zero mean,  $\sigma = 1.0$ ). Since the noise actually changes the data, it is not surprising to observe that the extracted regions have changed slightly. However, the segmentation changes significantly in the upper left corner of the image where regions are merged, and in the lower left corner where a region is split. Interestingly, this outcome is similar to the level 2 segmentation of Ohlander's algorithm (see fig. 9).

Filtering of images often has to be done to remove noise or to simulate scaling [14]. Linear smoothing, e.g. by gaussian convolution, leads to blurring of edges and therefore poses a problem for DRG and other region growing methods. For example, at region boundaries, smoothing induces a transition area that might be wrongly labeled as a new region. We experimented with a simple non-linear filtering method, namely median filtering. As fig. 8 shows (we used a small  $3 \times 3$  window), this kind of filtering tends to enlarge the stable regions and to smooth their boundaries. Some previously separated regions are merged. Also, some artifacts are introduced.

As a future research direction, we will work on the design of methods to simulate multiscale which cooperate well with region growing algorithms like DRG. [12] includes several new results in the area of linear and non-linear scale space that we want to build upon. Another non-linear scaling operator based on morphological dilation-erosion has been proposed in [6]. An examination of the interdependencies between scaling and the parameter settings of the segmentation algo-

gorithm should also provide better results.

The run time of DRG depends on the data. It is faster to segment images with large, homogeneous areas than to segment noisy images with only small homogeneous areas. We also observed that most coinciding regions that were eliminated were grown from seeds located in holes of already existing regions. That is, the overall run time can be reduced by disallowing such seeds.

There is no generally accepted methodology in the field of computer vision which elucidates how to evaluate segmentations algorithms (for short discussions of this topic, see [8],[10]). Comparing different segmentation algorithms with each other is difficult mainly because they differ in the properties they try to satisfy and in the image domain they are working in. Of all the algorithms the author is aware of, Ohlander's algorithm [9] comes closest to the proposed DRG algorithm because it also tries to derive the segmentation parameters dynamically from the image data rather than to operate with a priori thresholds, and it does not assume high-level knowledge or a specific image domain.

In the implementation of Ohlander's algorithm that we used, the maximum number of recursive splits could be limited by the user. Fig. 9 shows the segmentation we obtained by setting the maximum level to two, while fig. 10 shows the result for level three. One can see that Ohlander's algorithm extracts well the diagonal road across the upper left corner of the original image. Like with DRG, the extracted regions often have rugged boundaries and small holes. On the other hand, Ohlander's algorithm clearly tends to split regions without any consideration of spatial connectivity, leading to artifacts. The most difficult problem with Ohlander's algorithm, however, lies with the decision whether to keep on splitting a region. Ideally, the shape of the histograms suggest an answer to this question, but in noisy images this is often not the case. For example, we feel that the level three image in fig. 10 is already oversegmented.

## 5 Conclusion

We have presented a dynamic region growing algorithm that has some novel features that are important for an application in image indexing and retrieval. DRG allows to arrive fully automatically at a reasonable segmentation in noisy images without having to rely on a priori knowledge. The monitoring of the region growing steps provides strong clues that can be exploited for the segmentation process. First results show that stable regions can be successfully extracted. The important question whether the extracted regions

are *useful* can only be answered relative to the application and has to remain open at this point. It is also obvious that the no-knowledge constraint puts a limit on what DRG can achieve.

### Acknowledgements

I would like to thank B. Woodham and D. Lowe for their valuable suggestions. V. Tucakov provided the implementation of and insights into Ohlander's algorithm.

### References

- [1] R. Adams, L. Bischof, "Seeded Region Growing" *IEEE Trans. PAMI*, Vol.16, No.6, pp.641-647, June 1994.
- [2] J. Canny, "A Computational Approach to Edge Detection", *IEEE Trans. PAMI*, Vol.8, No.6, pp.679-698, Nov.1986.
- [3] C. Chu, J. Aggarwal, "The Integration of Image Segmentation Maps Using Region and Edge Information", *IEEE Trans. PAMI*, Vol.15, No.12, pp.1241-1252, Dec.1993.
- [4] R. Falah, P. Bolon, J. Cocquerez, "A Region-Region and Region-Edge Cooperative Approach of Image Segmentation", *ICIP-94*, 1994.
- [5] R. Haralick, L. Shapiro, "Image Segmentation Techniques" *Computer Vision, Graphics, and Image Processing*, Vol.29, pp.100-132, 1986.
- [6] P. Jackway, M. Deriche, "Scale-Space Properties of the Multiscale Morphological Dilation-Erosion", *IEEE Trans. PAMI*, Vol.18, No.1, pp.38-51, Jan.1996.
- [7] M. Kass, A. Witkin, D. Terzopoulos, "Snakes: Active Contour Models" *ICCV-87*, London 1987.
- [8] W. Ng, C. Lee, "Comment on Using the Uniformity Measure for Performance Measure in Image Segmentation", *IEEE Trans. PAMI*, Vol.18, No.9, pp.933-934, Sept.1996.
- [9] R. Ohlander, K. Price, D. Reddy, "Picture Segmentation Using a Recursive Region Splitting Method", *Computer Graphics and Image Processing*, Vol.8, pp.313-333, 1978.
- [10] N. Pal, S. Pal, "A Review on Image Segmentation Techniques", *Pattern Recognition*, Vol.26, No.9, pp.1277-1294, 1993.
- [11] R. Picard, T. Minka, "Vision Texture for Annotation", *Journal of Multimedia Systems*, Vol.3, pp.3-14, 1995.
- [12] B. ter Haar Romeny, "Geometry-Driven Diffusion in Computer Vision", Kluwer 1994.
- [13] R. Ronfard, "Region-Based Strategies for Active Contour Models", *Int. Journal of Computer Vision*, Vol.13, NO.2, pp.229-251, 1994.
- [14] A. Witkin, "Scale-Space Filtering", *IJCAI-83*, Karlsruhe 1983.
- [15] S. Zhu, A. Yuille, "Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation" *IEEE Trans. PAMI*, Vol.18, No.9, pp.884-900, Sept.1996.




Figure 3: Original aerial image




Figure 4: Segmentation, original image




Figure 5: Stable regions, original image




Figure 8: Stable regions, median filtered image




Figure 6: Stable regions, rotated image



Figure 9: Ohlander's algorithm, level 2



Figure 7: Stable regions, noisy image



Figure 10: Ohlander's algorithm, level 3