

Raster Based Region Growing

Donald G. Bailey
Image Analysis Unit
Massey University
Palmerston North

ABSTRACT

In some image segmentation applications, region growing is more appropriate than simple thresholding or edge detection. One example of this is the location of intensity peaks in an image, where each peak is defined by the valley between it and any adjacent peaks. Conventional region growing techniques require a seed point for each region, and the region is grown by adding adjacent pixels which match the segmentation criterion.

Some segmentation criteria are amenable to raster based region growing. Each pixel in the image is examined in turn. If it matches the inclusion criterion of a region associated with one of the pixels above it, it is allocated to that region and labelled accordingly. If not, a new region is started. Segmentation parameters for each region are held in a table indexed by the region label. In this way multiple regions are grown simultaneously. One disadvantage of this approach is that a single region may have multiple start pixels depending on the region shape. These individual regions eventually merge into a single region when they meet.

Raster based region growing is fast since it only requires a single pass through the image to perform the initial segmentation, with an additional pass to relabel adjacent regions which have merged.

INTRODUCTION

Image segmentation is the splitting of an image into its meaningful constituent parts, or regions which have a common property [1]. There are two broad categories of image segmentation techniques: those based on detecting edges or a discontinuity in some property between regions, and those based on classifying each pixel to a region according to some local property. Edge based techniques locate the transitions between regions and the regions are defined by the edges found. Region based techniques assign pixels to one region or another, and the edges are defined implicitly. Each of these can be further split into global and incremental methods, as shown in figure 1. Global methods perform the segmentation globally, for every pixel in the image. Incremental methods generally start with a seed pixel and then extend the edge or region by adding neighbouring pixels.

	Edge Based	Region Based
Global	Edge detection and linking	Thresholding
Incremental	Boundary Tracking	Region Growing

Figure 1: Categorisation of image segmentation techniques.

Thresholding and other global region based methods have the advantage of speed and simplicity but in some applications may not be appropriate. It is in these situations that region growing techniques may be most applicable.

In its simplest form, region growing is the joining of neighbouring pixels or groups of pixels which have similar properties into larger regions [2]. The process is started by selecting a seed pixel for the region. Neighbouring pixels are considered as candidates for inclusion in the growing region based on a similarity or inclusion test. Similar pixels are added to the region and their neighbours are examined in the same way. This process is continued until no more pixels are added to the region, and a new seed is selected to start the next region.

One of the problems with such a technique is the selection of the initial starting points [1]. Using different seed pixels may result in differences in the regions being detected. Often, some form of preprocessing is required to detect suitable starting points appropriate to the segmentation criterion being used [1].

RASTER IMPLEMENTATION

The need for a preprocessing step increases the execution time of the region growing method. In tasks with simple inclusion criteria or where the seed selection is not critical to the success of the region growing, the speed of the process may be improved by implementing it in a raster based manner.

Rather than grow one region at a time, several regions are grown in parallel as the image is scanned in raster fashion. Since we are working on several regions concurrently, each region is given a unique label. As the pixels are classified, they are labelled according to the region to which they are assigned.

Each pixel in the image is examined in turn. It is considered as a candidate for inclusion in the region associated with each of the neighbouring pixels above. (The pixels above will have already been allocated to a region since the previous rows have been processed). If the pixel satisfies the segmentation or inclusion criteria of one of the neighbouring regions, it is included within that region, otherwise it becomes the seed pixel of a new region and is labelled accordingly.

This avoids the problem of finding suitable seed points, since any pixel which does not fit into an existing neighbouring region automatically becomes a seed for a new region. One significant difference between this approach and conventional region growing is that the seed pixel will be on the edge of the final region, rather than somewhere near the centre. If the results of the particular region growing method depend significantly on the location of the seed pixel then the inclusion criteria may need to be modified to reduce this dependence.

One drawback of growing several regions simultaneously is that segmentation parameters for each region must be held separately. This may be accomplished by holding the information about each region in a table indexed by the region label. The specific information accumulated for each region depends on the segmentation or inclusion criteria used.

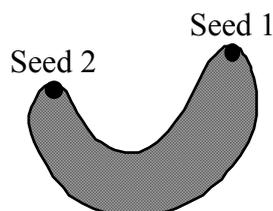


Figure 2: Multiple seed points for the same region.

One disadvantage of starting a new region before completing existing regions is that a single region may have multiple start pixels depending on the region shape. An example of such a region is shown in figure 2. This will result in single regions being split into several pieces unless some method of merging individual regions when they meet. Merging the regions may require an additional merge criterion which is not required with conventional region growing methods.

Another complication is that the regions which merge will have different labels. This may be overcome by having an additional merge label in the region parameter table. Before merging, the merge label points to the current region (as shown in the left of figure 3). Let us assume that regions 1 and 3 are the two regions being merged. Region data accumulated for the two regions is combined and the merge label is adjusted so that further references to region 3 are referred to entry 1 in the table. This avoids having to change all the pixels in the image which have been labelled as belonging to region 3. The relabelling can be performed after the complete image has been segmented, using the index and merge label as a lookup table.

Index (Label)	Region Data	Merge Label
1		1
2		2
3		3
4		4



Index (Label)	Region Data	Merge Label
1	Combined data	1
2		2
3	Invalid data	1
4		4

Figure 3: Region parameter table before and after merging regions 1 and 3.

In a sense, raster based region growing is a hybrid between the global and incremental methods described earlier, combining the incremental approach of adding pixels to existing regions with the global raster scanning and consideration of local information.

EXAMPLE - PEAK DETECTION

This process is best illustrated by example. One task where region growing is the most appropriate technique is the detection of intensity peaks in an image. Two applications which require peak detection are unsupervised multispectral classification (detecting peaks in a multi-dimensional intensity histogram) [1] and Hough transformation (detecting peaks corresponding to lines or circles in parameter space) [3].

In the input image, each intensity peak corresponds to a separate region, with the region boundaries defined by the positions of the valleys between the peaks. Simple thresholding is not appropriate, since the peaks are not all the same height and thresholding will only detect the top parts of the peaks (those greater than the threshold level).

With conventional region growing we can locate the pixels at the top of the peaks and use these as seed pixels. In practice, the most reliable way of achieving this is to locate the maximum pixel value within the image and use the position of that as the seed. During the growing process, a candidate pixel is added to the region if it is less than or equal to an adjacent pixel already belonging to the region. The region will therefore grow until a valley is reached, and the pixel values start increasing again. When no more pixels can be added, the image is then scanned to find the position of the pixel with the maximum value that has not yet been classified. This is used as the seed for the next region. This process is repeated until every pixel in the image has been classified.

The approach needs to be different for implementing peak detection in a raster based region growing algorithm since the seed pixels will not be the maximum values. In fact the seed pixels will be adjacent to valleys in the input image. This means that we need to grow from one valley, up to the peak, and then down the other side as far as the valley. Rather than working from top of the peaks down, the raster based technique works from the bottom up. The basic selection strategy works as follows.

For each pixel being examined, the maximum value of the 8 neighbouring pixels is determined. If the maximum is less than the value of the centre pixel then the top of the peak has been located. Otherwise the maximum pixel and the centre pixel are considered to belong to the same region. There are three cases corresponding to whether neither, one, or both of these pixels are labelled:

- if neither are labelled (for example if the maximum is below the centre pixel in the image) then a new region is started and a new label allocated to both the centre pixel and the maximum pixel;
- if one of the two pixels has already been labelled, then that label is used for the other pixel as well;
- if both have been labelled with the same label, then they already belong to the same peak. If the labels are different, then the two regions are merged and the region data table updated to reflect this.

In practice, the situation is a little more complicated. There may be several neighbouring pixels which share the maximum value. In this case, the maximum pixel with the greatest gradient is selected. If there are still several pixels, then one is chosen arbitrarily. A more serious difficulty occurs where the centre pixel value is the same as that of the maximum neighbour. There are three situations which may be represented in such a case:

- the pixels belong to a flat topped peak, in which case they belong to the same region;
- the pixels are part of a plateau part way up the side of a peak, and again they all belong to the same region;
- the pixels are part of a wide valley between two peaks, in which case some may have already been classified as belonging to each of the two peaks.

In the last case, we do not want the regions to be merged. To recognise this case, we need to be able to identify that we are in a valley. One way of achieving this is to record the current maximum value in each region being grown, and this information is used to modify the inclusion and merge criteria. To classify large regions of constant intensity correctly, when the centre pixel has the same value as the maximum neighbour, then all neighbouring pixels with this value are linked if they satisfy the merge criterion. The final link or merge strategy may be expressed by the following rules:

1. If the maximum neighbour is greater than the centre pixel value and
 - a) both are unlabelled, then a new region is started and the new label assigned to both the centre and the neighbour. The value of the neighbour is stored as the region maximum.
 - b) only the centre is unlabelled, then assign the label of the neighbour to the centre.
 - c) only the neighbour is unlabelled and if the centre value is the same as the region maximum then assign the label of the centre to the neighbour and update the region maximum to the value of the neighbour. This test with the region maximum prevents peaks on either side of wide flat region from being merged.
 - d) both pixels have the same label, then nothing further is required.
 - e) the pixels have different labels and if the centre value has the region maximum value then merge the two regions. Again, this test prevents merging of separate peaks adjacent to a wide flat area. Merging is achieved by the following steps:

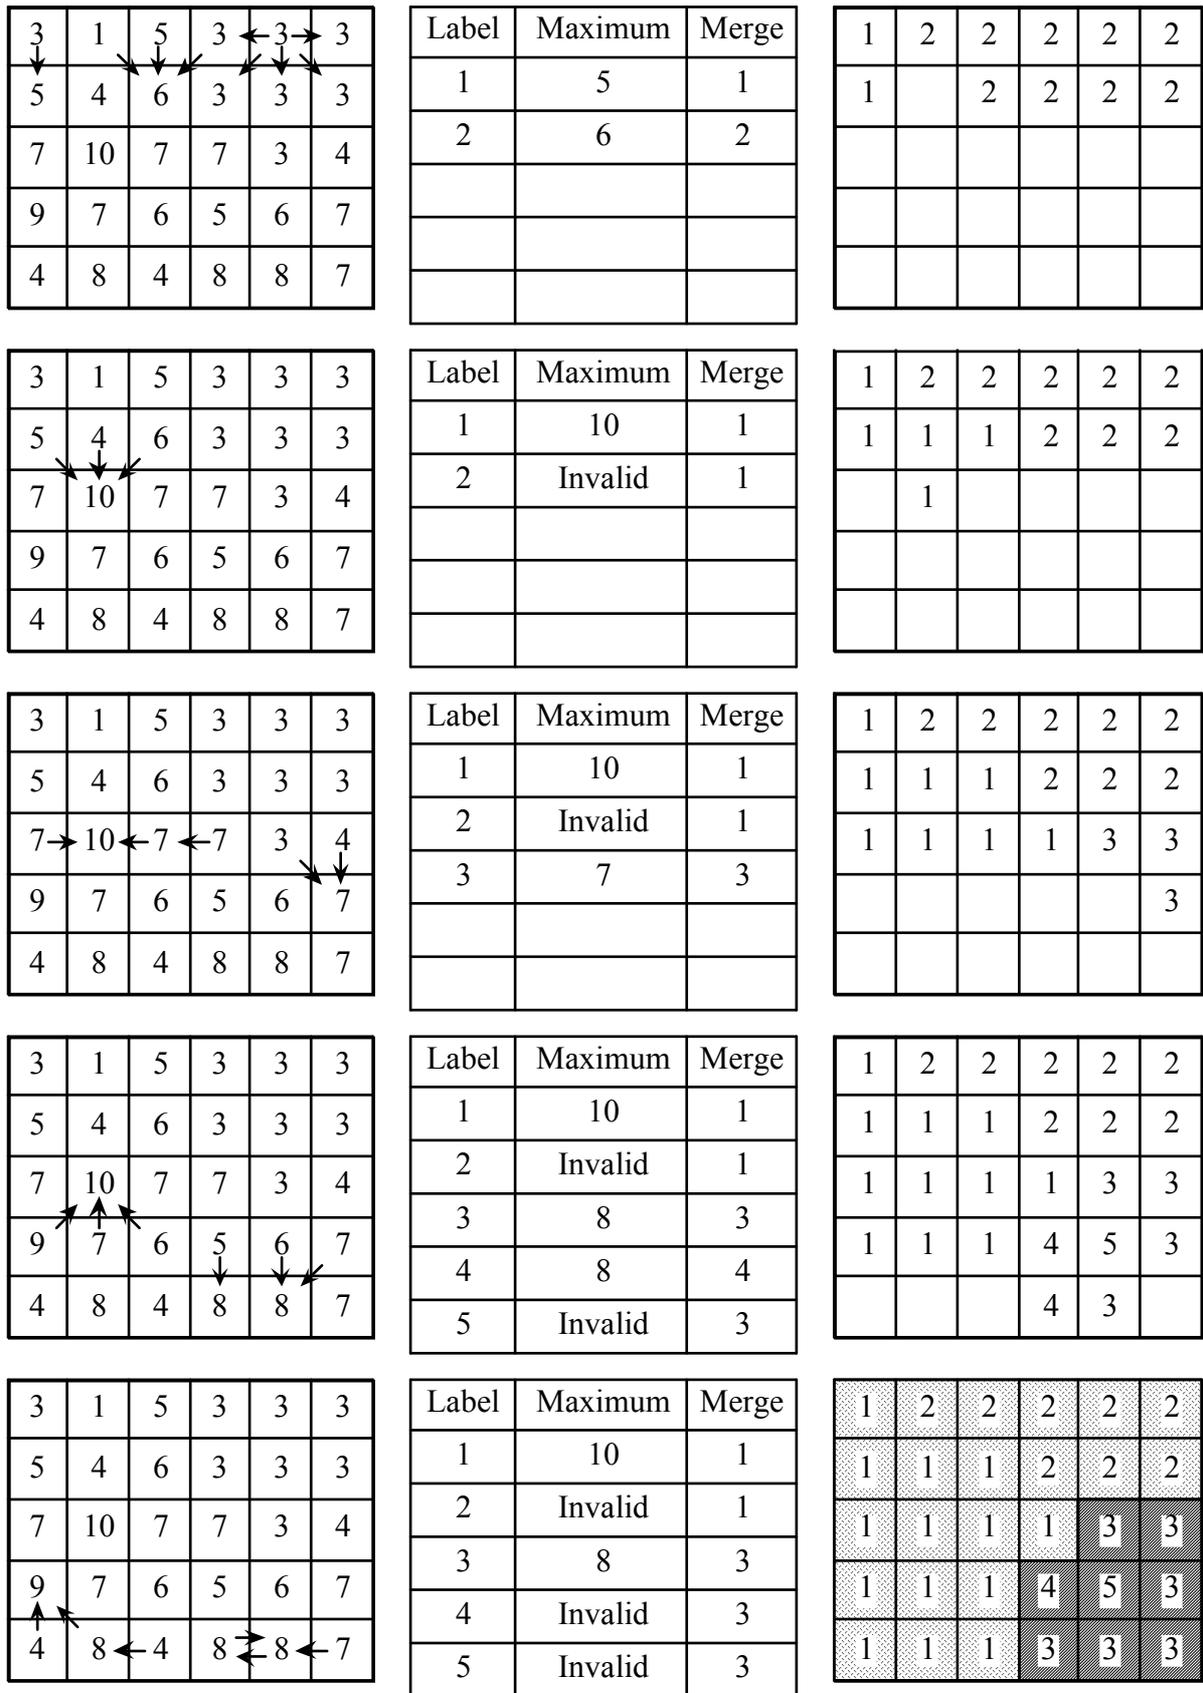


Figure 4: Raster based peak detection after processing each row of a sample image.
 The left column shows pixel linking; the centre column shows the region table;
 the right column shows the labelled image

- One of the labels is chosen as the new label and both pixels are assigned this label.
 - The region maximum of the new label is the greater of the region maximum of the regions being merged.
 - All references to the merged label in the region table are changed to the new label. This saves having to relabel the pixels in the image with the merged to the new label
2. If the maximum neighbour is less than the centre pixel value, then the peak has been found. If the centre pixel is unlabelled then a new region is started, the new label assigned only to the centre, and the centre value is stored as the region maximum.
 3. If the maximum neighbour has the same value as the centre pixel then all neighbours of the same value are considered.
 - a) If any of the neighbours being examined is labelled, then one of those labels is assigned to the centre. If none of the neighbours being considered is labelled then a new region is started and the new label assigned to the centre value.
 - b) All unlabelled neighbours under consideration are assigned the label of the centre value.
 - c) Any labelled neighbours with a label different from the centre pixel are merged only if the region maximum of either the neighbour or the centre is the same as the centre value.

A final relabelling pass through the image is required to relabel regions which have merged. The merge information from the region data table is used for this.

Figure 5 shows this process on a 5x6 image. In processing the first row, the group of 3s in the top right corner are detected and linked together. In the second row regions 1 and 2 merge and the right 3 pixels are not linked because they do not satisfy condition 1c) above. The third and fourth rows are reasonably straight forward. In the fifth row, the 8s are linked together, merging regions 3 and 4. The final labelling is shown by the shading

The above algorithm was implemented in VIPS [4] on a MicroVAX II. It took about 5 seconds to locate the peaks in a 128x128 image, about the same processing time as a median filter. A conventional region growing peak detection algorithm would require one pass through the image for each peak.

CONCLUSIONS

Region growing may be implemented in a raster based fashion if the inclusion criteria is simple and the seed selection is not critical. Raster based techniques grow several regions in parallel and only require a single pass through the image to perform the initial segmentation, with an additional pass to relabel adjacent regions which have merged. The region growing criteria may be more complex, however, since the seed pixels for each region are on the edge of the region, and multiple seeds may be detected for the same region.

REFERENCES

- [1] Gonzalez R.F. and Wintz P., *Digital Image Processing*, Addison Wesley, Reading Massachusetts, Second edition (1987).
- [2] Zucker S.W., "Region Growing: Childhood and Adolescence", *Computer Graphics and Image Processing*, vol 5, pp 382-399 (1976).
- [3] Williamson C.J., and Bailey D.G., "Automatic Graph Interpretation", 6th New Zealand Image Processing Workshop, Lower Hutt (1991).
- [4] Bailey D.G., and Hodgson R.M., "VIPS - a digital image processing algorithm development environment", *Image and Vision Computing*, vol 6, pp 176-184 (1986).