CIS Paper Review

Alexander Liu

Semi-Automatic Brain MRI Segmentation, Group 7

Motivation: The paper investigated was Hernandez, S. E., K. E. Barner, et al. (2005). "Region merging using homogeneity and edge integrity for watershed-based image segmentation." Optical Engineering **44(1): 017004-017014.** The premise of the paper was to discuss a hybrid method of the watershed transform for use in image segmentation. This paper was selected due to its relevance to our method of segmentation. We chose a watershed transform as well, and employed techniques similar to the ones used in this paper. However, we have only investigated the effects of using homogeneity methods, and therefore wanted to see the effects of combining this method with others. This paper also presented pre-processing methods, which we have also been investigating in an attempt to improve our results. Primarily, we wanted to investigate the effects of median filtering and edge integrity on segmentation results.

Summary: The watershed transform is a method in which the gradient magnitude image is obtained and used to interpret the image features as geographical surfaces. Pixels with highest gradient values are then identified as 'watershed lines', or lines that mark the boundaries of features in the image. Classical watershed methods yield numerous small regions (over-segmentation), and these regions must be merged in order to yield a meaningful result. In the past, linear and morphological filtering operations have been proposed to address this problem, but these methods often cause distortion of edge boundaries or are computationally intensive. Edge merge algorithms have also been proposed, and one commonly used is the region adjacency graph, which is an undirected graph whose nodes correspond to regions, and the edge weights are assigned according to a dissimilarity function. These functions can be based on homogeneity (versus the mean intensity) to find the weight of a given region versus its adjacent regions. Regions with the lowest dissimilarity via edge weight are then merged first for a desirable number of iterations. Other approaches base the edge weights on edge height, which is calculated by comparing gradient level along a common boundary. Borders with the smallest gradient are assigned the smallest weight, and merged first as well. While both these methods have been presented separately, they both, by themselves, each have drawbacks; homogeneity merging retains some undesirable background regions that constitute false boundaries, while edge weight merging results in many negligible regions. The paper proposes the combination of these two merging methods via dynamic weighting, in which proportions of both criterions are used to calculate the final weight of the edges. New pre-processing techniques are also discussed.

The overall flow of the transform used in the paper is as follows: Unfiltered Image -> Median Filtering -> Gradient Operator -> Median Filtering -> LIM and Morphological gradient calculation -> Gradient based region identification -> segmented image. In the first steps of a classical watershed transform, Gaussian filtering is employed. This removes plateaus (regions of uniform pixel value) through floating-point conversion, and is usually followed by a gradient operator, yielding a Gradient Magnitude Image (GMI):

$$\chi_{GMI} = \left\| \nabla\left(\mathcal{G}\left(\chi_{\mathrm{Orig}}\right)\right) \right\|$$

The overall result can be seen as a surface where the largest gradient magnitude values are boundaries of features in the image. The GMI in this paper was obtained using the Sobel operator, which is essentially convolution with a 3x3 kernel in either the x or the y direction, yielding approximate derivatives. At each point in the GMI, gradient approximations are combined with the following equations:

$$G_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}, G_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, G = \sqrt[2]{G_x^2 + G_y^2}$$

Following this, the image is interpreted using the largest gradient magnitude values as *Watershed Lines*, or lines that represent areas of largest contrast. The regions separated by these lines are *Catchment Basins*, and each catchment basin has a *Local Intensity Minimum*, or LIM, which is a place where gradient is lowest. This principle is illustrated in the figure below:

16.8	19.2	13.5	20.5	31.2	30.1
18.7	11.9	15.4	18.5	22.1	18.4
20.1	21.9	26.6	20.8	17.3	18.1
25.3	22.8	20.9	19.8	15.1	15.9
30.7	35.5	29.9	18.7	17.6	39.9
34.8	38.6	33.4	32.7	33.5	36.7

(a)

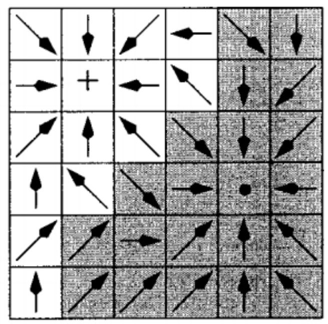

(b)

Each region has an LIM, and if you follow a given member of a region by its largest change in gradient, you will wind up at that region's LIM. In the modified method proposed by the authors, median pre-filtering is used instead of Gaussian filtering (Median filtering is simply taking the median of an nxn neighborhood for each pixel). The authors justify this by saying that this eliminates the false LIMs that were found in the previous methods. The authors then process this image with thresholding by taking the max between the image pixel intensity and a given threshold (in this case, they used a low threshold value to get rid of dark background island regions):

$$\chi_{GPP} = MAX\left[MED\left[\nabla\left(MED[\chi_{\text{Orig}}]\right)\right], \text{Threshold}\right]$$

In conclusion, the authors applied median filtering before and after applying the Sobel Operator, and then took the max of the resulting image with a threshold value to eliminate false background regions.

Following this, a Region Adjacency Graph (RAG) is constructed. Three different criterions of edge weight were investigated: edge orientation, region homogeneity, and edge integrity. The first, edge orientation, penalizes 135 degree edges (edges with such an orientation were merged first given their lower score). This was calculated by taking the regions and examining their border pixels and finding the proportion of pixels that are diagonal with respect to all border pixels. Regions with the largest ration of 'diagonal' pixels to total border pixels were merged first. The authors demonstrated the effectiveness of this algorithm at removing regions with diagonal borders, as shown in the space shuttle image:

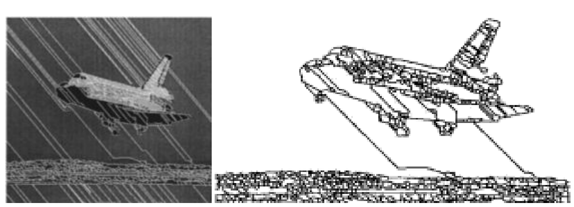

The next method, involves merging based on region homogeneity. The cost function for this is calculated as follows: Regions with the greatest similarity, in the mean sense, have the lowest edge weight. The dissimilarity function, which is shown below as δ^H, is defined as the product of the cardinalities of the two regions R_i and R_j multiplied by the square difference between their mean intensities normalized by the total number of pixels in the two regions combined.

$$\delta^H(R_i, R_j) = \frac{\|R_i\| \cdot \|R_j\|}{\|R_i\| + \|R_j\|}\left[\mu(R_i) - \mu(R_j)\right]^2$$

Following this, the authors investigated merging based on an edge integrity cost function. They calculated this cost by finding the max of two contour pixels from the gradient magnitude image. Any edge with pixels over a given threshold (which was the median of all height values) was considered a 'strong edge'. The proportion of strong pixels to border pixels was then again used to calculate the cost, shown below as δ^{ε}.

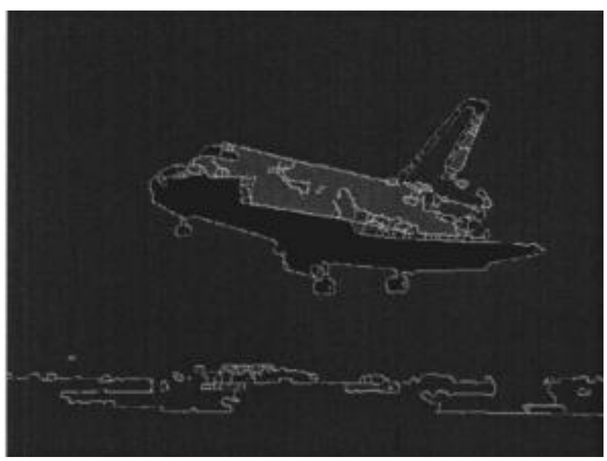

$$\delta^{\varepsilon}(R^i, R^j) = \frac{\|\varepsilon_S\|}{\|\varepsilon_B\|}$$

Using these two cost functions separately yielded sub-par results, and so the authors then went on to propose a dynamic weighting cost function, one that combines both region homogeneity and edge integrity. This was determined via the α parameter, which determined proportion of homogeneity edge weight. This was calculated in such a way so that alpha was the proportion of ‘small regions’ to total regions (since homogeneity merging was better for the smaller regions). A region was considered ‘small’ if its size fraction was less than a given threshold; that is, if the ratio of its pixels to the total pixels in the whole image was below a certain percentage (0.5%, in this case). The equations are enumerated below; W is the edge weight, R^H and R^E are homogeneity and Edge height values, and S_R is the set of all small regions in the RAG.

$$W = \alpha R^H + (1-\alpha) R^{\varepsilon}$$

$$S_R = \left\{ R^j : \frac{\|R^j\|}{\|\chi\|} < P_s \right\}$$

$$\alpha = \frac{\|S_R\|}{K}, 0 \le \alpha \le 1$$

The combined approach appeared to yield the most satisfactory results, as shown in the figure provided. The reduction of regions is outlined in a table, but the accuracy of the segmentation is not examined.

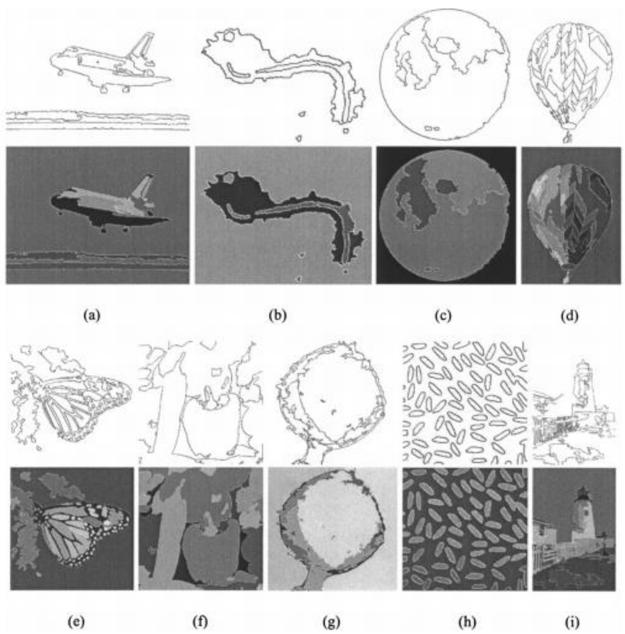

Image		Simulation parameters						
		Watershed procedure				Region-merging procedure		
		Pre-filter	Postfilter T_h	Postfilter Size	Initial regions	β	P_s (%)	Final regions
Name	Size							
Space Shuttle	423×559	3×3	20	3×3	934	0.7	0.5	25
Ebola	207×262	3×3	30	3×3	1413	0.7	0.5	9
Moon	400×400	3×3	30	3×3	1893	0.7	0.5	10
Balloon	512×366	3×3	30	3×3	754	0.7	0.5	100
Monarch	512×512	3×3	25	3×3	1747	0.7	0.5	100
Peppers	512×512	3×3	30	3×3	1697	0.7	0.5	50
Cell	482×502	3×3	60	3×3	2474	0.7	0.5	50
Rice	256×256	3×3	20	3×3	412	0.7	0.5	75
Lighthouse	768×512	3×3	40	3×3	2881	0.7	0.5	150

Critiques: The paper constantly tests the algorithm's efficacy using just a single space shuttle image as a test image. While this does give us a sense of how the algorithm performs on certain features (i.e. looking at how the algorithm segments certain border types, like the ones seen at the tail fin of the shuttle), it does not give us a very complete picture of how the algorithm performs on other images with various characteristics. While they did show the results on the test image sets, they did not show them in great detail, nor did they discuss the accuracy of the segmentations with respect to the true boundaries (the borders generated could have been overlaid upon the original images for inspection). Many times, they only investigated region reduction for many testing scenarios, opting not to show or verify whether or not the regions obtained were true to the actual contours. The algorithm also specifically targets 135 degree angle regions. While this does work well, and the 135 degree angles seen in the background of the shuttle image are indeed removed, it is unclear as to how common these regions are in other images besides the shuttle image. These regions could also become problematic in the case where the actual contours in the image are 135 degrees, and not just the ones in the background. The fact that they did not show the performance of this algorithm on other images also raises uncertainty as to the performance of the angled region reduction method in the general case, which is the case we are interested in due to the diversity of brain tumors, and specifically, Glioblastoma Multiforme, whose MRI scans are extremely diverse and sometimes noisy.

In addition to the image-specific techniques, we are also left with an ambiguous picture of the nature of the methods presented. The proposed methods were shown to have worked on the shuttle image, and possibly some of the other test images (again, these were not verifiable), but we are also unsure as to how these will perform on MRI images. Rather than show that the methods worked on a class of images, or even a set of images, the authors instead presented a case for an algorithm that worked on one image, and possibly for 8 other images. The results are mixed, with some minor improvements in spurious region elimination. Moreover, they are similar to ours, as the space shuttle was also shown to have a 'leak', in which some object regions (in this case, the space shuttle) become contiguous with the background due to poor border strength.

Overall, the paper presents and alternative method to the watershed transform that we see can work in some cases. While the paper does an excellent job of enumerating the techniques, it presents a surface investigation of the technique's efficacy and not an in-depth one. Rather than looking at the algorithm's actual results, the authors chose to instead extrapolate that the hybrid method's results kept the benefits of both while leaving behind their respective drawbacks. However, we can see that the hybrid method still yields failures, as seen in the space shuttle tail. We are unsure as to how this will perform with Brain MRI images, but it is highly unlikely that this hybrid method will address our leakage problem adequately. This paper shows us that our problem likely falls within the nature of the watershed transform itself.