

ECE661: Homework 6

Ahmed Mohamed (akaseb@purdue.edu)

October 28, 2014

1 Otsu Segmentation Algorithm

Given a grayscale image, my implementation of the Otsu algorithm follows these steps:

1. Construct a 256-level histogram h of the image, such that $h[i] = n_i$ is the number of pixels whose grayscale value equal to i .
2. Calculate the average grayscale value of the image.

$$\mu_T = \sum_1^L ip_i$$

where

$$p_i = n_i/N$$

and L is the total number of levels, and N is the total number of pixels in the image.

3. For each level in the histogram, calculate:
 - (a) The zeroth-order cumulative moment

$$\omega(k) = \sum_1^k p_i$$

- (b) The first-order cumulative moment

$$\mu(k) = \sum_1^k ip_i$$

- (c) The between-class variance

$$\sigma_B^2(k) = [\mu_T\omega(k) - \mu(k)]^2/[\omega(k)(1 - \omega(k))]$$

4. Choose $threshold = k^*$ such that $\sigma_B^2(k^*)$ is maximum.
5. Construct a mask whose pixels is 1 if the corresponding pixels in the original image is greater than the threshold, and 0 otherwise. This mask represents the foreground of the image.
6. Repeat this process for an arbitrary number of iterations. In each iteration, the new histogram will not contain the pixels below the threshold. The number of iterations is chosen manually to increase the quality of the results.

2 RGB Image Segmentation Using the Otsu algorithm

Given a color image, my implementation follows the following steps to extract the foreground of the image.

1. Separate the RGB color channels of the input image into three grayscale images.
2. Get the foreground mask for each channel using the Otsu algorithm as described earlier.
3. To merge the three masks together into a single foreground mask, we need to manually know about the colors of the foreground and the background of the image.
 - (a) For the lake image, the lake is mostly blue, while the background has other colors. Hence, we should treat the blue mask as the foreground mask, while the green and red masks as the background masks. Hence, the overall foreground mask is:

$$mask = mask_b \text{ AND } (NOT \ mask_g) \text{ AND } (NOT \ mask_r)$$

where $mask, mask_b, mask_g, mask_r$ are the overall, blue, green, and red masks respectively.

- (b) For the tiger image, there is no dominant color for the background or the foreground. Hence, we treat all the masks as foreground masks. Hence, the overall foreground mask is:

$$mask = mask_b \text{ AND } mask_g \text{ AND } mask_r$$

3 Texture-based Segmentation Using the Otsu algorithm

Given a color image, my implementation follows the following steps to extract the foreground of the image.

1. Convert the image into a grayscale image.
2. Create a grayscale image whose pixels represent the variance of the grayscale values of the $N * N$ window around the corresponding pixels in the original grayscale image.
3. Do the previous steps for $N = 3, N = 5,$ and $N = 7$ to get three grayscale images that represent the texture-based features of the original image. These three grayscale images are considered three channels of an image.
4. Get the foreground mask for each channel using the Otsu algorithm as described earlier.
5. To merge the three masks together into a single foreground mask, we need to manually know about the structure of the foreground and the background of the image.
 - (a) For the lake image, the pixels of the lake have almost no variance. Hence, we treat the three masks as background masks. Hence, the overall foreground mask is:

$$mask = (NOT \ mask_1) \text{ AND } (NOT \ mask_2) \text{ AND } (NOT \ mask_3)$$

where $mask, mask_1, mask_2, mask_3$ are the overall mask and the masks of the individual channels respectively.

- (b) For the tiger image, the pixels of the tiger have larger variance than others. Hence, we treat the three masks as foreground masks. Hence, the overall foreground mask is:

$$mask = mask_1 \text{ AND } mask_2 \text{ AND } mask_3$$

where $mask, mask_1, mask_2, mask_3$ are the overall mask and the masks of the individual channels respectively.

4 Noise Elimination

The foreground masks may be noisy. To eliminate the noise we use combinations of dilation and erosion as follows:

1. Erosion then Dilation with a square window to remove the noise in the background.
2. Dilation then Erosion with a square window to remove the noise in the foreground.

For the lake image, we needed both approaches to eliminate the noise in both the foreground and the background. However, for the tiger image, we used only the second approach because there were many holes in the foreground, and applying the first approach would have removed most of the mask.

5 Contour Extraction

Given a binary mask, the contour is the foreground pixels that touch the background. Thus, we follow the following steps:

1. For each pixel in the binary mask:
 - (a) If the pixel value is 0, it doesn't belong to the contour.
 - (b) If the pixel value is 1, and all adjacent pixels in a 3x3 window (8-connectivity) are 1, it doesn't belong to the contour.
 - (c) Only if the pixel is 1, and one or more of its adjacent pixels are zero, it is considered on the contour.

6 More Observations

1. The results quality of the texture-based segmentation is much better than the image-based segmentation.
2. Texture-based segmentation consumes more time than the image-based segmentation.
3. Image-based segmentation require human knowledge of the different colors of foregrounds and backgrounds. For example, my implementation for the lake image will not produce good results for other images, because we had to use the information that the foreground is blue. This might not be the case with other images.
4. Texture-based segmentation requires less knowledge about the foreground and the background. It will divide the image into areas with high variances, and areas with low variances. We still need to take a human decision about which is the foreground and which is the background.
5. In brief, The Otsu algorithm is fast, but very limited. It doesn't produce very good results. Using texture-based segmentation instead of image-based segmentation increases the quality of the results, especially with challenging images such as the tiger image.
6. Erosion and Dilation are very effective removing the noise in both the background and the foreground. However, one should be careful choosing the window size, because this process could remove the entire mask.

7 Results

7.1 Image 1: Lake



Figure 1: The input image

7.1.1 RGB Segmentation




Figure 2: The three color channels (BGR) of the image




Figure 3: The foreground mask using the blue channel of the image, and the corresponding foreground




Figure 4: The foreground masks using the green channel of the image (4 iterations of Otsu's algorithm are used.)




Figure 5: The foreground of the image corresponding to the masks of the green channel (4 iterations of Otsu's algorithm are used.)




Figure 6: The foreground masks using the red channel of the image (3 iterations of Otsu's algorithm are used.)




Figure 7: The foreground of the image corresponding to the masks of the red channel (3 iterations of Otsu's algorithm are used.)




Figure 8: The overall mask using the masks of the three channels, and the corresponding foreground




Figure 9: The final overall mask after removing the noise, and the corresponding foreground




Figure 10: The final contour of the image

7.1.2 Texture-based Segmentation




Figure 11: The three channels of the image representing the texture-based features (corresponding to the windows 3x3, 5x5, and 7x7 respectively)




Figure 12: The foreground mask using the first channel of the image, and the corresponding foreground




Figure 13: The foreground mask using the second channel of the image, and the corresponding foreground




Figure 14: The foreground mask using the third channel of the image, and the corresponding foreground




Figure 15: The overall mask using the masks of the three channels, and the corresponding foreground



Figure 16: The final overall mask after removing the noise, and the corresponding foreground

7.2 Image 2: Tiger



Figure 17: The input image

7.2.1 RGB Segmentation



Figure 18: The three color channels (BGR) of the image



Figure 19: The foreground mask using the blue channel of the image, and the corresponding foreground



Figure 20: The foreground mask using the green channel of the image, and the corresponding foreground




Figure 21: The foreground mask using the red channel of the image, and the corresponding foreground




Figure 22: The overall mask using the masks of the three channels, and the corresponding foreground




Figure 23: The final overall mask after removing the noise, and the corresponding foreground




Figure 24: The final contour of the image

7.2.2 Texture-based Segmentation




Figure 25: The three channels of the image representing the texture-based features (corresponding to the windows 3x3, 5x5, and 7x7 respectively)



Figure 26: The foreground mask using the first channel of the image, and the corresponding foreground




Figure 27: The foreground mask using the second channel of the image, and the corresponding foreground




Figure 28: The foreground mask using the third channel of the image, and the corresponding foreground




Figure 29: The overall mask using the masks of the three channels, and the corresponding foreground



Figure 30: The final overall mask after removing the noise, and the corresponding foreground



Figure 31: The final contour of the image

8 Source Code

The following is the entire Python source code.

```

1 import cv2
2 import numpy as np
3
4
5 # The input image file name.
6 FILE_NAME = 'images/1.jpg'
7
8 # Whether to use texture features or not (RGB values).

```

```

9 USE_TEXTURE = False
10
11
12 def main():
13
14     # Read the input image.
15     image = cv2.imread(FILE_NAME)
16     cv2.imshow(FILE_NAME, image)
17
18     if not USE_TEXTURE:
19         s = ''
20
21         # Segment using Otsu's algorithm on the image.
22         mask = otsu_rgb(image, inverted_masks=[0, 1, 1],
23                        iterations=[1, 4, 3], s=s)
24     else:
25         s = '_t'
26
27         # Build an image with the texture-based features of the grayscale image.
28         texture_image = get_texture_image(image)
29
30         # Segment using Otsu's algorithm on the texture-based image.
31         mask = otsu_rgb(texture_image, inverted_masks=[1, 1, 1],
32                        iterations=[1, 1, 1], org_image=image, s=s)
33
34         # Erosion then Dilation to remove the noises in the background.
35         kernel = np.ones((17, 17), np.uint8)
36         mask = cv2.erode(mask, kernel)
37         mask = cv2.dilate(mask, kernel)
38
39         # Dilation the Erosion to remove the noises in the foreground.
40         kernel = np.ones((5, 5), np.uint8)
41         mask = cv2.dilate(mask, kernel)
42         mask = cv2.erode(mask, kernel)
43
44         cv2.imshow(inject(FILE_NAME, 'filtered_mask{}'.format(s)), mask)
45         cv2.imwrite(inject(FILE_NAME, 'filtered_mask{}'.format(s)), mask)
46         cv2.imshow(inject(FILE_NAME, 'filtered_foreground{}'.format(s)),
47                   cv2.bitwise_and(image, image, mask=mask))
48         cv2.imwrite(inject(FILE_NAME, 'filtered_foreground{}'.format(s)),
49                   cv2.bitwise_and(image, image, mask=mask))
50
51         # Extract the contour
52         contour = extract_contour(mask)
53         cv2.imshow(inject(FILE_NAME, 'contour{}'.format(s)), contour)
54         cv2.imwrite(inject(FILE_NAME, 'contour{}'.format(s)), contour)
55
56         while not cv2.waitKey(50) & 0xFF == 27: pass
57         cv2.destroyAllWindows()
58
59
60 def inject(image_name, suffix):
61     return '{}-{}.{}'.format(image_name.split('.')[0], suffix,
62                              image_name.split('.')[1])
63
64
65 def otsu_rgb(image, inverted_masks=[0, 0, 0],
66             iterations=[1, 1, 1], org_image=None, s=''):
67
68     # The original image is just used for the results.
69     if org_image is None:
70         org_image = image
71
72     # Initialize the overall mask.
73     overall_mask = np.ndarray((image.shape[0], image.shape[1]), np.uint8)
74     overall_mask.fill(255)
75

```

```

76 # For each channel in the three color channels:
77 for c in xrange(3):
78
79     print 'Processing_channel:_' , c
80
81     # The image representing channel c.
82     channel_image = np.zeros_like(image)
83     channel_image[:, :, c] = image[:, :, c]
84     cv2.imshow(inject(FILE_NAME, 'channel_{{}}'.format(c, s)),
85                channel_image)
86     cv2.imwrite(inject(FILE_NAME, 'channel_{{}}'.format(c, s)),
87                channel_image)
88
89     # The mask of channel c.
90     mask = None
91
92     # For an arbitrary number of iterations: perform the segmentation
93     # using Otsu's algorithm.
94     for i in xrange(iterations[c]):
95         # Perform the segmentation using Otsu's algorithm.
96         mask = otsu(image[:, :, c], mask)
97
98         # Save the results.
99         cv2.imshow(inject(FILE_NAME, 'mask_{{}}_{{}}'.format(c, i, s)), mask)
100        cv2.imwrite(inject(FILE_NAME, 'mask_{{}}_{{}}'.format(c, i, s)), mask)
101        cv2.imshow(inject(FILE_NAME, 'foreground_{{}}_{{}}'.format(c, i, s)),
102                   cv2.bitwise_and(org_image, org_image, mask=mask))
103        cv2.imwrite(inject(FILE_NAME, 'foreground_{{}}_{{}}'.format(c, i, s)),
104                   cv2.bitwise_and(org_image, org_image, mask=mask))
105
106        # Calculate the overall mask as the logical and of masks after inverting
107        # the masks that are indicated in inverted_masks.
108        if inverted_masks[c] == 1:
109            overall_mask = cv2.bitwise_and(overall_mask, cv2.bitwise_not(mask))
110        else:
111            overall_mask = cv2.bitwise_and(overall_mask, mask)
112
113        # Save the results.
114        cv2.imshow(inject(FILE_NAME, 'mask_{{}}'.format(s)), overall_mask)
115        cv2.imwrite(inject(FILE_NAME, 'mask_{{}}'.format(s)), overall_mask)
116        cv2.imshow(inject(FILE_NAME, 'foreground_{{}}'.format(s)),
117                   cv2.bitwise_and(org_image, org_image, mask=overall_mask))
118        cv2.imwrite(inject(FILE_NAME, 'foreground_{{}}'.format(s)),
119                   cv2.bitwise_and(org_image, org_image, mask=overall_mask))
120
121    return overall_mask
122
123
124 def otsu(image, mask=None):
125
126     # The histogram of grayscale levels.
127     histogram = [0] * 256
128
129     # The total number of pixels in the mask.
130     pixels_num = 0
131
132     # The average grayscale value for the entire image (masked by the mask).
133     mu_t = 0
134
135     # Initialize the histogram based on the pixels of the image.
136     for r in xrange(image.shape[0]):
137         for c in xrange(image.shape[1]):
138             if mask is None or mask[r][c] != 0:
139                 pixels_num += 1
140                 mu_t += image[r][c]
141                 histogram[image[r][c]] += 1
142

```

```

143     # The average grayscale value for the entire image (masked by the mask).
144     mu_t = float(mu_t) / pixels_num
145
146     # The cumulative probability of pixels less than or equal level i.
147     omega_i = 0
148
149     # The cumulative average grayscale value less than or equal level i.
150     mu_i = 0
151
152     # The final chosen threshold.
153     threshold = -1
154
155     # The maximum  $\sigma_b^2$  corresponding to the final chosen threshold.
156     max_sigma_b = -1
157
158     # For every grayscale level in the histogram:
159     for i in xrange(256):
160
161         # The number of pixels in the grayscale level i.
162         n_i = histogram[i]
163
164         # The probability of pixels in level i.
165         p_i = n_i / float(pixels_num)
166
167         # Update the cumulative probability of pixels less than or equal
168         # level i, and the cumulative average grayscale value less than or
169         # equal level i.
170         omega_i += p_i
171         mu_i += i * p_i
172
173         # Ignore the very first levels and the very last levels that don't
174         # contain any pixels. For these levels,  $\sigma_{b_i}$  will cause division
175         # by zero exception.
176         if omega_i == 0 or omega_i == 1:
177             continue
178
179         # Update the between-class variance  $\sigma_b^2$ .
180         sigma_b_i = (mu_t * omega_i - mu_i) ** 2 / (omega_i * (1 - omega_i))
181
182         # Compare the between-class variance  $\sigma_b^2$  to the maximum,
183         # and update the best threshold.
184         if sigma_b_i > max_sigma_b:
185             threshold = i
186             max_sigma_b = sigma_b_i
187
188     print threshold
189
190     # The image of the output mask.
191     output_mask = np.zeros_like(image)
192
193     if threshold == -1:
194         return output_mask
195
196     # For each pixel in the input image:
197     for r in xrange(image.shape[0]):
198         for c in xrange(image.shape[1]):
199             # Set the corresponding output mask pixel to 1 if the pixel
200             # values is greater than the threshold.
201             if image[r][c] > threshold:
202                 output_mask[r][c] = 255
203
204     return output_mask
205
206
207 def get_texture_image(color_image):
208
209     # The grayscale version of the image.

```



```

210 image = cv2.cvtColor(color_image, cv2.COLOR_BGR2GRAY)
211
212 # The texture-based image.
213 texture_image = np.zeros_like(color_image)
214
215 # The different window sizes used for the texture features.
216 window_size = [3, 5, 7]
217
218 # For each different window size:
219 for i, w in enumerate(window_size):
220
221     # Half of the window size.
222     d = w / 2
223
224     # For each pixel in the input image:
225     for r in xrange(d, image.shape[0] - d):
226         for c in xrange(d, image.shape[1] - d):
227             # Calculate the variance of the pixel values in the window,
228             # and store it in the texture_image.
229             texture_image[r][c][i] = np.int(
230                 np.var(image[r - d: r + d + 1, c - d: c + d + 1]))
231
232     return texture_image
233
234
235 def extract_contour(image):
236
237     # The contour image.
238     contour = np.zeros_like(image)
239
240     # For each pixel in the input image:
241     for r in xrange(1, image.shape[0] - 1):
242         for c in xrange(1, image.shape[1] - 1):
243
244             # If the pixel is non-zero and there exist a zero pixel in
245             # the surrounding 3x3 window, the pixel is considered on the
246             # contour.
247             if image[r][c] != 0 and \
248                 np.min(image[r - 1: r + 2, c - 1: c + 2]) == 0:
249                 contour[r][c] = 255
250
251     return contour
252
253
254 if __name__ == "__main__":
255     main()

```

Listing 1: The entire Python source code