

ECE661: Homework 7

Ahmed Mohamed (akaseb@purdue.edu)

November 6, 2014

1 Character Recognition

In order to recognize the characters in a set of test images using a training image, my implementation follows the following steps for each of the images:

1. Segment the foreground using Otsu's algorithm as shown in Section 2. The foreground represents the letters on the image. The output of this process is a mask with non-zero pixels corresponding to the foreground pixels.
2. Perform component labeling as shown in Section 3 in order to label the different disconnected components in the foreground mask. The output of this process is an image whose pixel values are labels (integers). Each unique label correspond to a component (which is actually a letter) in the image.
3. Clean the components as shown in Section 4 by removing the very big components which represent the background, not the actual characters, as well as the very small components (likely to be noise).
4. Perform Harris corner detection to find the sharpest N (I use $N = 9$) corners in each component as shown in Section 5.
5. Construct a shape vector for each single component (character) in the image as shown in Section 6.
6. If this is the training image, keep the shape vectors for the matching process. If this is a test image, find the best component match of the shape vector of each component in the test image as shown in Section 7.

2 RGB Image Segmentation Using the Otsu algorithm

Given a color image, my implementation follows the following steps to extract the foreground of an 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 in the next subsection.
3. To merge the three masks together into a single foreground mask, we observe that the foreground is always colorful (high pixel values), and the background is either black or white. We also note that letters have different colors. So we need the foreground to be the union of all the foregrounds from the three RGB channels. Hence, the overall foreground mask is:

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

where $mask$, $mask_b$, $mask_g$, $mask_r$ are the overall, blue, green, and red masks respectively. The training image is a little different since the foreground is black and the background is white, so the masks have to be inverted.

2.1 Grayscale Otsu Segmentation

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.

3 Component Labeling

Given the foreground mask, the output of this process is a labels image whose pixel values are labels (integers). Each unique label correspond to a component (which is actually a letter) in the image. My component labeling implementation follows the following steps:

1. First pass: assign temporary labels for connected pixels, and record labels equivalences. To achieve that, for each pixel in the mask:
 - (a) If the pixel value is 0, assign its label to 0.
 - (b) Construct the equivalence set of labels for this pixel as the labels of the neighboring pixels (west, north west, north, north east) whose pixel value is equal to the value of this pixel.

- (c) Note that, in case of 4-connectivity, we only check for the west and north pixels. We needed the 4-connectivity with the first test image because two characters were 8-connected.
 - (d) Check the size of the equivalence set:
 - i. If the size of the equivalence set is equal to 0, this means that none of the neighboring pixels has labels, so we assign a new label to this pixel.
 - ii. If the size of the equivalence set is equal to 1, this means that only one neighboring pixel has a label, so we assign this label to this pixel.
 - iii. If the size of the equivalence set is more than 1, this means that the neighboring pixels have different labels. In this case we need to record this equivalence. So, we add this equivalence set to a global list of equivalence lists. And we choose the least label from the equivalence set to this pixel.
 - iv. Note that when adding the equivalence set to the list of equivalence sets, the equivalence set may share some labels with any of the previous equivalence sets. In this case, we merge all the sets who intersect into a single equivalence set.
2. Second pass: resolve equivalences. To achieve that, for each pixel with non-zero label in the labels image:
- (a) Check the list of equivalence sets. If the label exists in one of the sets, assign the least label from the set.
 - (b) If not, leave the label as is.
 - (c) Note that after this pass, all the labels who exist in the same equivalence set will have the least label in the set.

4 Cleaning Components

This process has two main goals.

1. Remove the components whose size is more than 30,000 pixels. These components are likely to be the background, because some images have white backgrounds. This is done by creating a histogram that represents the frequency of each label. After that, we iterate over the image to set the labels whose frequency is more than 30,000 to the label 0 (background).
2. Remove the components whose size is less than 100 pixels. These components are likely to be noise.
3. For convenience and implementation ease, the method replace the non-continuous labels with continuous ones. The output of the component labeling could contain labels 2, 20, 40, etc. After this step, the new labels will be 1, 2, 3, etc.

5 Harris Corner Detector

My Harris corner detection program follows the following steps to find a maximum of N corners in each component in an image given a specific scale σ .

1. Initialize a list of corners for each label, and a list of the ratios of the corresponding corners. These lists categorizes the corners by their labels, and will be used to choose the sharpest corners.
2. Smooth the input image by applying a Gaussian filter with the given σ .

3. Find the x-derivative and y-derivative of each pixel in the smoothed image by applying a Haar wavelet filter. The Haar window size is the least even number that is greater than 4σ . If $\sigma = 1.2$, *window size* = 6. As a result, the following operator will be used to find dx :

$$\begin{pmatrix} -1 & -1 & -1 & 1 & 1 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 \end{pmatrix}$$

And the following operator will be used to find dy :

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 & -1 \end{pmatrix}$$

In order to perform this convolution efficiently, I use the integral image of the smoothed image.

4. For each labeled pixel in the image:

- (a) Construct the following matrix using the $5\sigma \times 5\sigma$ window around the pixel:

$$C = \begin{pmatrix} \sum d_x^2 & \sum d_x d_y \\ \sum d_x d_y & \sum d_y^2 \end{pmatrix}$$

- (b) If the pixel is not a corner, one of the eigen values will be very small. At each pixel, the corner strength (how likely the pixel is a corner) is given by the following relation:

$$CornerStrength = \lambda_1 \lambda_2 - k(\lambda_1 + \lambda_2)^2$$

Where k is a constant equal to 0.04. To avoid finding the eigen values of the matrix, it is known that:

$$Det(C) = \lambda_1 \lambda_2$$

$$Tr(C) = \lambda_1 + \lambda_2$$

So the corner strength can now be computed as:

$$CornerStrength = Det(C) - kTr(C)^2$$

- (c) The pixel is considered a corner if its Corner Strength is greater than a certain threshold `HARRIS_THRESHOLD` which I set to 10^{12} .
- (d) If the pixel is considered a corner, add the pixel to the list of corners of its label, and the ratio to the list of ratios of its label.
5. A Non-maximum Suppression process is performed. In other words, a corner is eliminated if one of its neighbors has a larger corner strength.
6. Use the lists of corners and ratios to choose the sharpest N corners for each label, i.e. the corners with the largest ratios.

6 Constructing Shape Vectors

To construct a shape vector that represent a component (character), my method follows the following steps for each component.

1. Calculate the center of the component. The center is computed as follows:

$$x_{center} = (x_{max} + x_{min})/2$$

$$y_{center} = (y_{max} + y_{min})/2$$

where x_{center} and y_{center} are the x and y coordinates of the center point. x_{max} and y_{max} x_{min} and y_{min} are the coordinates of the extreme points in the two directions.

2. Calculate the angle of each corner with the center of the component:

$$\tan^{-1}((y_{corner} - y_{center})/(x_{corner} - x_{center}))$$

The angle is then adjusted if it is negative.

3. Calculate the angle between each two consecutive corners by sorting the angles and then subtracting each two consecutive angles. The angle between the corners is the same as the arc length between the projections of the corners on the unit circle. The angle is easier to implement than the arc length between the projected corners, so I didn't project the corners.

$$ArcLength = r\theta$$

where r is the radius of the circle (1 because it is unit circle), and θ is the angle between the two corners.

4. The shape vector is the vector of the angles between the lines joining the component center and the corners.
5. As an enhancement, I calculate the distance between each corner and the center pixel. That's because some letters have corners with the same angles, but the distances to the corners are different. Then I append the list of normalized distances to the list of normalized angles in the shape vector. This enhancement will be evaluated later.
6. Optimally, we would have N angles between the N corners. But for some components, there is no N corners, so we have to pad the shape vector with zeros.

7 Matching Letters

Given a shape vector of a component, to find the best training component match, my method follows the following procedures:

1. For each element in the shape vector:
 - (a) In order to make the matching process rotation invariant, Circularly rotate the vector such that this element is the first.
 - (b) For each training shape vector:
 - i. Find the Euclidean distance between the rotated shape vector and the training shape vector.
 - ii. If the Euclidean distance is the least one so far, the best component match is the component corresponding to the current training shape vector.

8 More Thoughts

The results section show that the overall average recognition accuracy for all the letters for all the test images is 31%. Here are my observations:

1. In my opinion, Harris corner detection is the main reason of the low performance due to the following reasons:
 - (a) It is hard to tune the variable N (the maximum number of corners for each letter. If I use large N, Harris will include many not-sharp corners in the letters with few corners (e.g. I, O), which makes the matching process hard. If I use small N, Many important corners won't be included from the letters with many corners (e.g. M, W).
 - (b) It is hard to set a Harris threshold that is suitable for all the images. If the threshold is low, this will include many bad corners. If it is high, this will miss many important corners.
 - (c) In case of a letter with many corners (e.g. M, W, A), not all the corners will be included. For example, the letter A has 11 corner, but the maximum I use is 9, so two of the corners won't be included each time. However, the 9 sharpest corners differ from image to image, which result in a shape vector that is greatly different from the training shape vector.
2. The fonts used for different images are not the same, which makes it very hard to recognize the letters. For example, the letter O has no corners in the training image, but it has many edges in the last test image (in the word "Hollywood". This shows that the approach of using the corners to recognize letters will always be limited. This is very clear because the statistics indicate that the images with the best recognition accuracy are images 3 and 4 because their fonts are similar to the training set, while no letters recognised in Image 2 because the font is very different.

As an enhancement, I try using a modified shape vector that includes the distance to the center point. I expected that the results would be better if the shape vector encodes the normalized distance to the corner as well. That's because some letters have corners with the same angles, but the distances to the corners are different. Using this enhancement, the accuracy for some images has greatly increased (image 2 from 0% to 15%). However, the overall accuracy for all the images was almost the same 30% (vs. 31% without the enhancement).

Many other things can be done to achieve better results such as using larger training set. The training set should contain many fonts and many ways to write each letter, so that it is easier to recognize letters from the test set.

9 Results

9.1 Training Image




Figure 1: The input training image




Figure 2: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.




Figure 3: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.




Figure 4: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$




Figure 5: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)

9.2 Test Image 1



Figure 6: The input test image



Figure 7: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.



Figure 8: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.



Figure 9: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$



Figure 10: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)



Figure 11: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.3 Test Image 2



Figure 12: The input test image



Figure 13: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.



Figure 14: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.



Figure 15: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$



Figure 16: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)



Figure 17: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.4 Test Image 3




Figure 18: The input test image




Figure 19: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.




Figure 20: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.

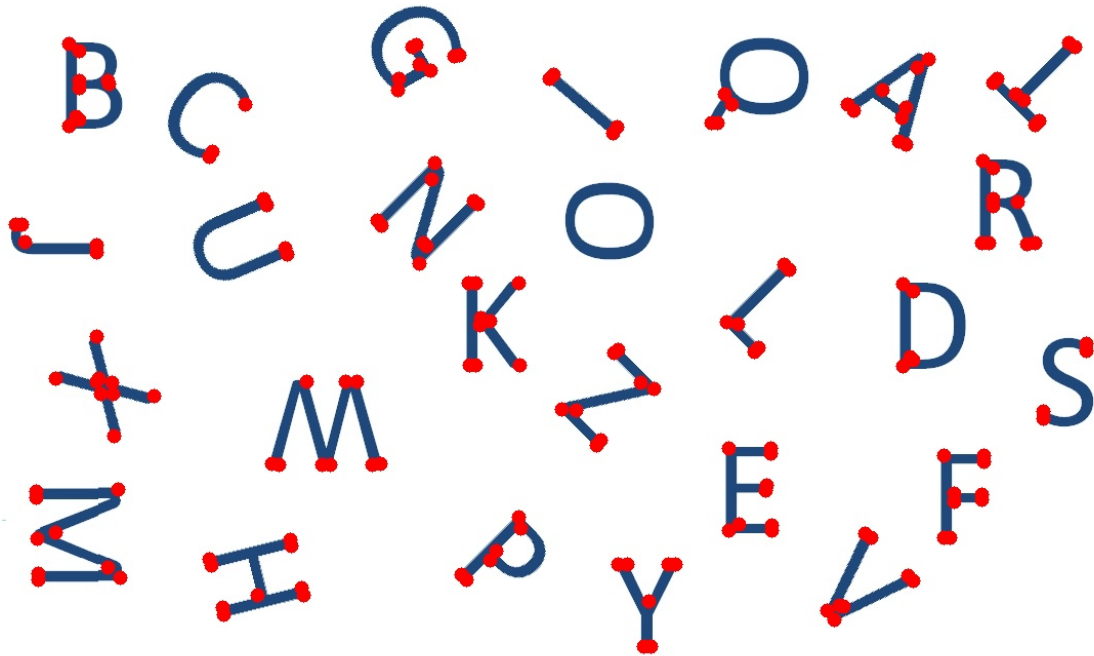


Figure 21: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$

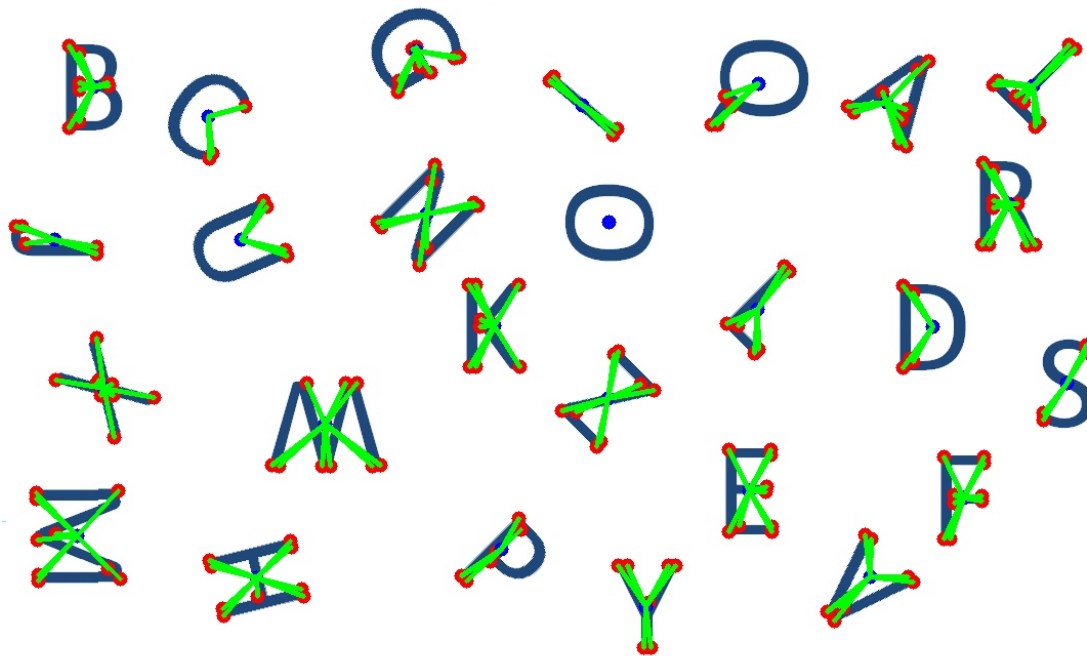


Figure 22: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)

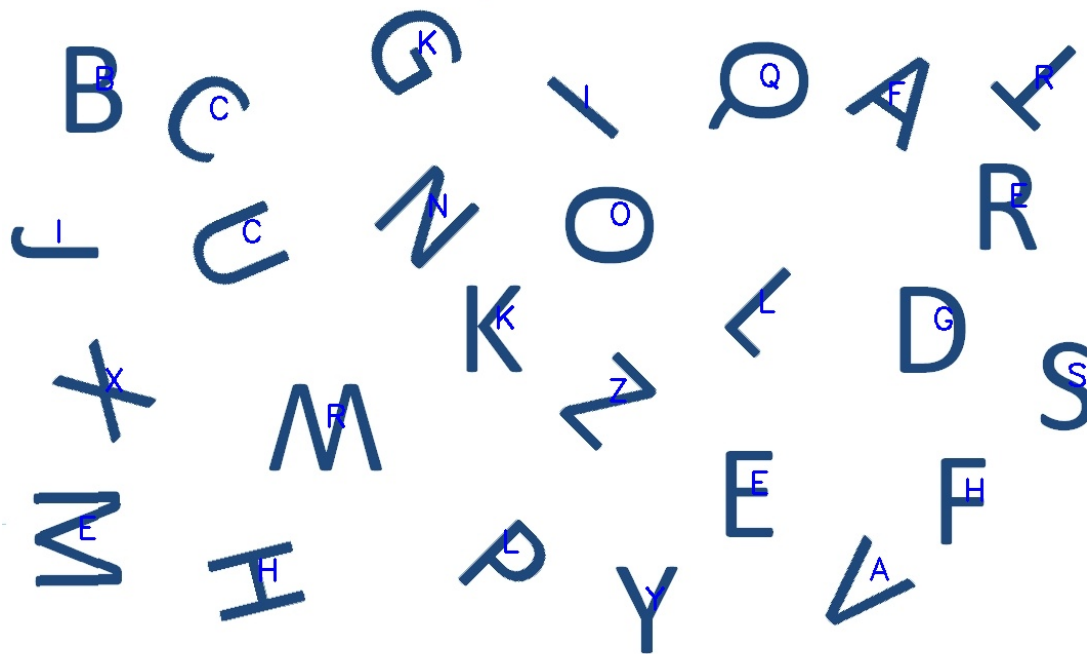


Figure 23: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.5 Test Image 4

**ΑΓΓ THE
ΜΟΒΓΔ IS A
STAGE**

**ALL THE WEN
AND MOWEN
WEBEΓΛ
PLAYERS**

Figure 24: The input test image

**ALL THE
MONEY IS A
STAGE**

**ALL THE MEN
AND WOMEN
WEBER
PLAYERS**

Figure 25: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.



Figure 26: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.

VTT THE
MOBTD IS A
STAGE

ALL THE MEN
AND WOMEN
WEBEΓA
PLAYERS

Figure 27: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$

VFF THE
MOBFD IS A
STAGE

ALL THE MEN
AND WOMEN
WEBER
PLAYERS

Figure 28: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)

VTT THE
MOBTD IS A
STAGE
ALL THE WEN
AND MOWEN
WEBEY
PLAYERS

Figure 29: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.6 Test Image 5






Figure 31: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.



Figure 32: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.



Figure 33: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$




Figure 34: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)



Figure 35: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.7 Test Image 6



Figure 36: The input test image



Figure 37: Otsu: the foreground mask of the image. White pixels are the foreground, while black pixels are the background.



Figure 38: Component labeling: an image representing the output of the component labeling process. Each unique label (component) is given a random color for visualization. Note that this image is after the component cleaning process as well.




Figure 39: Harris corners (red circles) for each component. Note that some characters don't have enough corners. The number of corners per character will not exceed $N = 9$




Figure 40: The shape vector of each component is the N-vector of angles between each successive pair of the N (green) lines that join the component center (blue) with the corners (red)



Figure 41: The output of the character recognition process. The overlaid blue characters are the output of the recognition process for the corresponding underlaid characters.

9.8 Statistics

	Image 1		Image 2		Image 3		Image 4		Image 5		Image 6		Overall		
	#	C	#	C	#	C	#	C	#	C	#	C	#	C	%
A	1	0	1	0	1	0	6	4	1	0	0	0	10	4	40.00%
B	1	0	1	0	1	1	0	0	1	0	0	0	4	1	25.00%
C	1	0	1	0	1	1	0	0	1	0	0	0	4	1	25.00%
D	1	0	1	0	1	0	2	2	1	1	1	0	7	3	42.86%
E	1	0	1	0	1	1	8	3	1	0	0	0	12	4	33.33%
F	1	0	1	0	1	0	0	0	1	0	0	0	4	0	0.00%
G	1	0	1	0	1	0	1	1	1	0	0	0	5	1	20.00%
H	1	0	1	0	1	1	2	1	1	1	1	0	7	3	42.86%
I	1	1	1	0	1	1	1	1	1	0	0	0	5	3	60.00%
J	1	0	1	0	1	0	0	0	1	0	0	0	4	0	0.00%
K	1	0	1	0	1	1	0	0	1	0	0	0	4	1	25.00%
L	1	0	1	0	1	1	7	0	1	1	2	1	13	3	23.08%
M	1	1	1	0	1	0	3	0	1	0	0	0	7	1	14.29%
N	1	0	1	0	1	1	3	1	1	1	0	0	7	3	42.86%
O	1	1	1	0	1	1	2	2	1	1	3	0	9	5	55.56%
P	1	0	1	0	1	0	1	1	1	1	0	0	5	2	40.00%
Q	1	0	1	0	1	1	0	0	1	1	0	0	4	2	50.00%
R	1	1	1	0	1	0	3	1	1	0	0	0	7	2	28.57%
S	1	0	1	0	1	1	3	0	1	0	0	0	7	1	14.29%
T	1	0	1	0	1	0	3	2	1	1	0	0	7	3	42.86%
U	1	0	1	0	1	0	0	0	1	0	0	0	4	0	0.00%
V	1	1	1	0	1	0	0	0	1	0	0	0	4	1	25.00%
W	1	0	1	0	1	0	2	0	1	0	1	0	7	0	0.00%
X	1	1	1	0	1	1	0	0	1	0	0	0	4	2	50.00%
Y	1	0	1	0	1	1	2	1	1	0	1	0	7	2	28.57%
Z	1	0	1	0	1	1	0	0	1	1	0	0	4	2	50.00%
Tot	26	6	26	0	26	14	49	20	26	9	9	1	162	50	
%		23.08%		0.00%		53.85%		40.82%		34.62%		11.11%		30.86%	

Figure 42: Statistics of the recognition Process WITHOUT the enhancement of the corners distances. The first column contains the letters. The second and third columns contain the actual number of characters in Image 1 and the correctly recognized letters respectively. Columns 4-13 are the same for all the other test images. Column 14 and 15 contain the actual number of characters in all the images and the correctly recognized letters respectively. Column 16 contains the overall recognition accuracy for each letter. The last two rows show the overall statistics for all letters in each image. The bold-ed number shows that the overall average recognition accuracy is 31correctly recognized out of 162 letters)

	Image 1		Image 2		Image 3		Image 4		Image 5		Image 6		Overall		
	#	C	#	C	#	C	#	C	#	C	#	C	#	C	%
A	1		1	1	1	1	6		1	1	0		10	3	30.00%
B	1		1		1	1	0		1		0		4	1	25.00%
C	1		1		1		0		1		0		4	0	0.00%
D	1		1		1		2	2	1	1	1		7	3	42.86%
E	1		1		1	1	8		1		0		12	1	8.33%
F	1		1		1		0		1		0		4	0	0.00%
G	1		1		1		1		1		0		5	0	0.00%
H	1		1		1	1	2	2	1	1	1		7	4	57.14%
I	1	1	1		1	1	1	1	1		0		5	3	60.00%
J	1		1	1	1		0		1		0		4	1	25.00%
K	1		1		1		0		1		0		4	0	0.00%
L	1	1	1		1	1	7	4	1	1	2	1	13	8	61.54%
M	1		1		1	1	3		1		0		7	1	14.29%
N	1		1		1		3		1		0		7	0	0.00%
O	1	1	1		1	1	2	2	1	1	3		9	5	55.56%
P	1		1		1		1	1	1		0		5	1	20.00%
Q	1		1	1	1		0		1	1	0		4	2	50.00%
R	1		1		1	1	3	1	1		0		7	2	28.57%
S	1		1		1		3	3	1		0		7	3	42.86%
T	1		1		1		3	2	1		0		7	2	28.57%
U	1	1	1		1		0		1		0		4	1	25.00%
V	1		1		1		0		1		0		4	0	0.00%
W	1		1	1	1	1	2	1	1		1		7	3	42.86%
X	1		1		1	1	0		1		0		4	1	25.00%
Y	1		1		1		2	1	1		1	1	7	2	28.57%
Z	1		1		1		0		1	1	0		4	1	25.00%
Tot	26	4	26	4	26	11	49	20	26	7	9	2	162	48	
%	15.38%		15.38%		42.31%		40.82%		26.92%		22.22%		29.63%		

Figure 43: Statistics of the recognition Process WITH the enhancement of the corners distance. The first column contains the letters. The second and third columns contain the actual number of characters in Image 1 and the correctly recognized letters respectively. Columns 4-13 are the same for all the other test images. Column 14 and 15 contain the actual number of characters in all the images and the correctly recognized letters respectively. Column 16 contains the overall recognition accuracy for each letter. The last two rows show the overall statistics for all letters in each image. The bold-ed number shows that the overall average recognition accuracy is 31 (correctly recognized out of 162 letters)

10 Source Code

The following is the entire Python source code.

```
1 import sys
2 import cv2
3 import numpy as np
4 from math import atan2, pi
5
6 # Harris Corner Detection Threshold
7 HARRIS_THRESHOLD = 1e12
8
9 # The number of corners in one letter.
10 CORNERS_NUM = 9
11
12
13 def main():
14
15     # The shape vectors of the training image.
16     training_vectors = None
17
18     # The order of the letters in the training image according to the labels
19     # numbers.
20     training_letters = 'CGABDEFHILJKLMNOQSPRTUVWXYZ'
21
22     for i in xrange(7):
23
24         print 'Image', i
25
26         # The file name of the input image.
27         file_name = 'images/{}.jpg'.format(i)
28
29         # Read the input image.
30         image = cv2.imread(file_name)
31
32         image_clone = image.copy()
33
34         # Segment using Otsu's algorithm on the image.
35         if i != 6:
36             mask = otsu_rgb(image, file_name, inverted_masks=[1, 1, 1],
37                             is_and=0)
38         else:
39             mask = otsu_rgb(image, file_name, inverted_masks=[0, 0, 0])
40
41         # Perform component labeling. Use four connectivity only for the
42         # testing image number 1.
43         labels_image = label_components(mask, four_connectivity=(i == 1))
44
45         # Clean the components by removing the very large components.
46         clean_components(labels_image)
47
48         # Visualize the labels by assigning a random color to each label.
49         show_labels(labels_image, file_name)
50
51         # Convert the color image to grayscale.
52         gray_image = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
53
54         # Find the corners in the grayscale image using Harris corner detection.
55         components_corners = harris(gray_image, labels_image)
56
57         # Mark the corners in the first image.
58         for component in components_corners:
59             for corner in component:
60                 cv2.circle(image, (corner[0], corner[1]), 7, (0, 0, 255), -1)
61
62         # Save the results.
63         cv2.imshow(inject(file_name, 'corners'), image)
64         cv2.imwrite(inject(file_name, 'corners'), image)
```



```

65
66     # Get the shape vectors.
67     components_centers, components_vectors = get_shape_vectors(
68         labels_image, components_corners)
69
70     # The shape vectors of the first image, are the training vectors.
71     if i == 0:
72         training_vectors = components_vectors
73
74     for j, center in enumerate(components_centers):
75         cv2.circle(image, (center[0], center[1]), 7, (255, 0, 0), -1)
76
77         for corner in components_corners[j]:
78             cv2.line(image, (center[0], center[1]), (corner[0], corner[1]),
79                 (0, 255, 0), 3)
80
81     # Save the results.
82     cv2.imshow(inject(file_name, 'features'), image)
83     cv2.imwrite(inject(file_name, 'features'), image)
84
85     if i == 0:
86         # Print the letters on the image.
87         for j, center in enumerate(components_centers):
88             cv2.putText(image_clone, training_letters[j],
89                 center, cv2.FONT_HERSHEY_SIMPLEX, 1, (255, 0, 0), 2)
90     else:
91         # Match the letters in the image with the training image.
92         matchings = recognize_components(training_vectors,
93             components_vectors)
94         # Print the letters on the image.
95         for j, center in enumerate(components_centers):
96             cv2.putText(image_clone, training_letters[matchings[j]],
97                 center, cv2.FONT_HERSHEY_SIMPLEX, 1, (255, 0, 0), 2)
98
99     # Save the results.
100    cv2.imshow(inject(file_name, 'recognition'), image_clone)
101    cv2.imwrite(inject(file_name, 'recognition'), image_clone)
102
103    while not cv2.waitKey(50) & 0xFF == 27:
104        pass
105    cv2.destroyAllWindows()
106
107
108    def inject(image_name, suffix):
109        return '{}-{}.{}'.format(image_name.split('.')[0], suffix,
110            image_name.split('.')[1])
111
112
113    def otsu_rgb(image, file_name, inverted_masks=[0, 0, 0], is_and=1,
114        iterations=[1, 1, 1], org_image=None, s=''):
115
116        # The original image is just used for the results.
117        if org_image is None:
118            org_image = image
119
120        # Initialize the overall mask.
121        overall_mask = np.zeros((image.shape[0], image.shape[1]), np.uint8)
122        if is_and:
123            overall_mask.fill(255)
124
125        # For each channel in the three color channels:
126        for c in xrange(3):
127
128            # The image representing channel c.
129            channel_image = np.zeros_like(image)
130            channel_image[:, :, c] = image[:, :, c]
131

```

```

132     # The mask of channel c.
133     mask = None
134
135     # For an arbitrary number of iterations: perform the segmentation
136     # using Otsu's algorithm.
137     for i in xrange(iterations[c]):
138         # Perform the segmentation using Otsu's algorithm.
139         mask = otsu(image[:, :, c], mask)
140
141         """
142         # Save the results.
143         cv2.imwrite(inject(file_name, 'mask_{}_{}'.format(c, i, s)), mask)
144         """
145
146     # Invert the masks that are indicated in inverted_masks.
147     if inverted_masks[c] == 1:
148         mask = cv2.bitwise_not(mask)
149
150     # Calculate the overall mask as the logical and/or of masks.
151     if is_and:
152         overall_mask = cv2.bitwise_and(overall_mask, mask)
153     else:
154         overall_mask = cv2.bitwise_or(overall_mask, mask)
155
156     # Save the results.
157     cv2.imwrite(inject(file_name, 'mask{}'.format(s)), overall_mask)
158
159     return overall_mask
160
161
162 def otsu(image, mask=None):
163
164     # The histogram of grayscale levels.
165     histogram = [0] * 256
166
167     # The total number of pixels in the mask.
168     pixels_num = 0
169
170     # The average grayscale value for the entire image (masked by the mask).
171     mu_t = 0
172
173     # Initialize the histogram based on the pixels of the image.
174     for r in xrange(image.shape[0]):
175         for c in xrange(image.shape[1]):
176             if mask is None or mask[r][c] != 0:
177                 pixels_num += 1
178                 mu_t += image[r][c]
179                 histogram[image[r][c]] += 1
180
181     # The average grayscale value for the entire image (masked by the mask).
182     mu_t = float(mu_t) / pixels_num
183
184     # The cumulative probability of pixels less than or equal level i.
185     omega_i = 0
186
187     # The cumulative average grayscale value less than or equal level i.
188     mu_i = 0
189
190     # The final chosen threshold.
191     threshold = -1
192
193     # The maximum  $\sigma_b^2$  corresponding to the final chosen threshold.
194     max_sigma_b = -1
195
196     # For every grayscale level in the histogram:
197     for i in xrange(256):
198

```

```

199     # The number of pixels in the grayscale level i.
200     n_i = histogram[i]
201
202     # The probability of pixels in level i.
203     p_i = n_i / float(pixels_num)
204
205     # Update the cumulative probability of pixels less than or equal
206     # level i, and the cumulative average grayscale value less than or
207     # equal level i.
208     omega_i += p_i
209     mu_i += i * p_i
210
211     # Ignore the very first levels and the very last levels that don't
212     # contain any pixels. For these levels, sigma_b_i will cause division
213     # by zero exception.
214     if omega_i == 0 or omega_i == 1:
215         continue
216
217     # Update the between-class variance  $\sigma_b^2$ .
218     sigma_b_i = (mu_t * omega_i - mu_i) ** 2 / (omega_i * (1 - omega_i))
219
220     # Compare the between-class variance  $\sigma_b^2$  to the maximum,
221     # and update the best threshold.
222     if sigma_b_i > max_sigma_b:
223         threshold = i
224         max_sigma_b = sigma_b_i
225
226     # The image of the output mask.
227     output_mask = np.zeros_like(image)
228
229     if threshold == -1:
230         return output_mask
231
232     # For each pixel in the input image:
233     for r in xrange(image.shape[0]):
234         for c in xrange(image.shape[1]):
235             # Set the corresponding output mask pixel to 1 if the pixel
236             # values is greater than the threshold.
237             if image[r][c] > threshold:
238                 output_mask[r][c] = 255
239
240     return output_mask
241
242
243 def label_components(image, four_connectivity=False):
244
245     # The output image of labels.
246     labels_image = np.zeros(image.shape, np.uint16)
247
248     # The next label to use.
249     next_label = 1
250
251     # The list of labels equivalence sets.
252     equivalence_sets = []
253
254     # The first pass: assign temporary labels, and record equivalences.
255     # For each pixel in the image:
256     for r in xrange(image.shape[0]):
257         for c in xrange(image.shape[1]):
258
259             # Process only non-zero pixels.
260             if image[r][c] == 0:
261                 continue
262
263             # The equivalence set of labels for this pixel.
264             equivalence_set = set([])
265

```

```

266     # If this pixel value is equal to the west pixel:
267     if c - 1 >= 0 and image[r][c] == image[r][c - 1]:
268         # Add the label to the equivalence set.
269         equivalence_set.add(labels_image[r][c - 1])
270
271     # If this pixel value is equal to the north pixel:
272     if r - 1 >= 0 and image[r][c] == image[r - 1][c]:
273         # Add the label to the equivalence set.
274         equivalence_set.add(labels_image[r - 1][c])
275
276     # If it is 8-connectivity, check the north west and north east.
277     if not four_connectivity:
278         # If this pixel value is equal to the north west pixel:
279         if (r - 1 >= 0 and c - 1 >= 0 and
280             image[r][c] == image[r - 1][c - 1]):
281             # Add the label to the equivalence set.
282             equivalence_set.add(labels_image[r - 1][c - 1])
283
284         # If this pixel value is equal to the north east pixel:
285         if (r - 1 >= 0 and c + 1 < image.shape[1] and
286             image[r][c] == image[r - 1][c + 1]):
287             # Add the label to the equivalence set.
288             equivalence_set.add(labels_image[r - 1][c + 1])
289
290     # Check the number of labels in the equivalence set.
291     if len(equivalence_set) == 0:
292         # If no labels in the equivalence set, assign new label.
293         label = next_label
294         next_label += 1
295     elif len(equivalence_set) == 1:
296         # If only one label in the equivalence set, choose it.
297         label = min(equivalence_set)
298     else:
299         # Choose the least label for the current pixel.
300         label = min(equivalence_set)
301         # If more than one label in the equivalence set:
302         # For every set in the global list of equivalence sets.
303         for i in xrange(len(equivalence_sets) - 1, -1, -1):
304
305             # If the current pixel equivalence set share any labels
306             # with the equivalence set in the list:
307             if bool(equivalence_set & equivalence_sets[i]):
308                 # Merge the two into the pixel equivalence set.
309                 equivalence_set |= equivalence_sets[i]
310                 # Delete the merged set from the list of sets.
311                 del equivalence_sets[i]
312
313             # Add the pixel equivalence set to the list.
314             equivalence_sets.append(equivalence_set)
315
316     labels_image[r][c] = label
317
318     # The second pass: resolve equivalences.
319     # For each pixel in the image.
320     for r in xrange(labels_image.shape[0]):
321         for c in xrange(labels_image.shape[1]):
322
323             # Process only non-zero labels.
324             if labels_image[r][c] != 0:
325
326                 # The new label to be assigned after resolving equivalences.
327                 new_label = -1
328
329                 # If the pixel label belongs to one of the equivalence sets,
330                 # assign the least label from the equivalence set.
331                 for equivalence_set in equivalence_sets:
332                     if labels_image[r][c] in equivalence_set:

```

```

333         new_label = min(equivalence_set)
334         break
335
336         # Assign the final label in case of equivalences, otherwise,
337         # leave the label as is.
338         if new_label != -1:
339             labels_image[r][c] = new_label
340
341     return labels_image
342
343
344 def clean_components(labels_image):
345
346     # This method removes the components whose size is larger than 30000
347     # pixel. These components are likely to be the background.
348
349     # The frequencies of the labels in the image.
350     labels_frequencies = np.bincount(labels_image.ravel())
351
352     # The set of labels to remove.
353     removed_labels = set([])
354
355     # For each label, if its frequency exceeds 30000, add it to the set.
356     for i in xrange(1, len(labels_frequencies)):
357         if labels_frequencies[i] > 30000 or labels_frequencies[i] < 100:
358             removed_labels.add(i)
359
360     # For each pixel in the image.
361     for r in xrange(labels_image.shape[0]):
362         for c in xrange(labels_image.shape[1]):
363
364             # If the label is one of the labels to remove, remove it.
365             if labels_image[r][c] != 0 and labels_image[r][c] in removed_labels:
366                 labels_image[r][c] = 0
367
368     # Then, the method assigned new labels so that they are continuous 1, 2, ...
369
370     # The current unique labels (not continuous).
371     labels = np.unique(labels_image)
372
373     # For each pixel in the image.
374     for r in xrange(labels_image.shape[0]):
375         for c in xrange(labels_image.shape[1]):
376
377             # Assign the index of the label. The indices are continuous.
378             if labels_image[r][c] != 0:
379                 labels_image[r][c] = np.where(labels == labels_image[r][c])[0]
380
381
382 def show_labels(labels_image, file_name):
383
384     # The image with the colored components
385     labels_colors = np.zeros((labels_image.shape[0], labels_image.shape[1], 3),
386                             np.uint8)
387
388     # Get the unique labels.
389     labels = np.unique(labels_image)
390     print 'Components-#', len(labels) - 1
391
392     # Generate a list of random colors for the list of unique labels.
393     random_colors = np.random.random_integers(255, size=(len(labels), 3))
394
395     # For each pixel in the image.
396     for r in xrange(1, labels_image.shape[0]):
397         for c in xrange(1, labels_image.shape[1]):
398             # Process only non-zero labels.
399             if labels_image[r][c] != 0:

```

```

400         # Color the label with the corresponding color.
401         label_index = np.where(labels == labels_image[r][c])[0]
402         labels_colors[r][c][:] = random_colors[label_index]
403
404     # Save the results.
405     cv2.imwrite(inject(file_name, 'labels'), labels_colors)
406     cv2.imshow(inject(file_name, 'labels'), labels_colors)
407
408
409 def harris(image, labels_image, sigma=1.2):
410
411     # The unique labels.
412     labels = np.unique(labels_image)
413
414     # The array for the corners of each component.
415     components_corners = [[] for i in xrange(len(labels) - 1)]
416     # The array of the corresponding ratios.
417     corners_ratios = [[] for i in xrange(len(labels) - 1)]
418
419     smoothed_image = cv2.GaussianBlur(image, (0, 0), sigma)
420
421     integral_image = np.zeros(image.shape)
422     for y in range(image.shape[0]):
423         integral_image[y][0] = smoothed_image[y][0]
424         for x in range(1, image.shape[1]):
425             integral_image[y][x] = (smoothed_image[y][x] +
426                                     integral_image[y][x - 1])
427
428     for x in range(image.shape[1]):
429         for y in range(1, image.shape[0]):
430             integral_image[y][x] += integral_image[y - 1][x]
431
432     haar_window = int(np.ceil(sigma * 4))
433     haar_window += haar_window % 2
434
435     dx = np.zeros(image.shape)
436     dy = np.zeros(image.shape)
437     for y in range(haar_window / 2, image.shape[0] - haar_window / 2):
438         for x in range(haar_window / 2, image.shape[1] - haar_window / 2):
439             dx[y][x] = get_dx(integral_image, x, y, haar_window)
440             dy[y][x] = get_dy(integral_image, x, y, haar_window)
441
442     harris_window = int(np.ceil(sigma * 5))
443     harris_window -= (1 - harris_window % 2)
444
445     corners = []
446     ratio_image = np.zeros_like(dx)
447
448     # For each pixel in the input image:
449     for y in range(harris_window / 2, image.shape[0] - harris_window / 2):
450         for x in range(harris_window / 2, image.shape[1] - harris_window / 2):
451
452             # Don't process un-labeled pixels.
453             if labels_image[y][x] == 0:
454                 continue
455
456             # The summation of the sqr(x-derivative), sqr(y-derivative),
457             # and x-derivative * y-derivative over the Harris window.
458             dx2 = 0
459             dxdy = 0
460             dy2 = 0
461
462             # For each pixel in the Harris window around the pixel (x,y).
463             for j in range(y - harris_window / 2,
464                             y + harris_window / 2 + 1):
465                 for i in range(x - harris_window / 2,
466                                 x + harris_window / 2 + 1):

```

```

467
468         # Get the x-derivative and y-derivative of the pixel (i,j).
469         dx_ij = dx[j][i]
470         dy_ij = dy[j][i]
471
472         # Add to the summation of the sqr(x-derivative),
473         # sqr(y-derivative), and x-derivative * y-derivative over
474         # the Harris window.
475         dx2 += np.square(dx_ij)
476         dxdy += dx_ij * dy_ij
477         dy2 += np.square(dy_ij)
478
479     if dx2 == 0 and dy2 == 0 and dxdy == 0:
480         continue
481
482     dx2 /= harris_window
483     dy2 /= harris_window
484     dxdy /= harris_window
485
486     # Compute the determinant and the trace of the C matrix at the
487     # pixel (x,y)
488     det_c = dx2 * dy2 - np.square(dxdy)
489     tr_c = dx2 + dy2
490
491     # Compute the ratio between the determinant and the trace squared.
492     #ratio = det_c / np.square(tr_c)
493     ratio = det_c - 0.04 * np.square(tr_c)
494     ratio_image[y][x] = ratio
495
496     if ratio >= HARRIS.THRESHOLD:
497         components_corners[labels_image[y][x] - 1].append((x, y))
498         corners_ratios[labels_image[y][x] - 1].append(ratio)
499
500     """
501     # If the ratio is above the HARRIS.THRESHOLD, the pixel is
502     # considered a corner.
503     if ratio >= HARRIS.THRESHOLD:
504         corners.append((x, y))
505     """
506
507
508 for i in xrange(len(components_corners)):
509     for j in xrange(len(components_corners[i]) - 1, -1, -1):
510
511         if not is_corner(components_corners[i][j][0],
512                          components_corners[i][j][1],
513                          ratio_image, harris_window):
514             del components_corners[i][j]
515             del corners_ratios[i][j]
516
517 for i in xrange(len(components_corners)):
518     if len(components_corners[i]) < CORNERS.NUM:
519         print len(components_corners[i])
520
521     max_indices = sorted(range(len(corners_ratios[i])),
522                          key=lambda x: corners_ratios[i][x])[-CORNERS.NUM:])
523     components_corners[i] = [components_corners[i][j] for j in max_indices]
524
525 return components_corners
526
527
528 def fast_convolute(integral_image, x1, y1, x2, y2):
529     return integral_image[y2][x2] - integral_image[y2][x1 - 1] - \
530            integral_image[y1 - 1][x2] + integral_image[y1 - 1][x1 - 1]
531
532
533 def get_dx(integral_image, x, y, window_size):

```

```

534
535     half1 = fast_convolute(integral_image ,
536                           x, y - window_size / 2,
537                           x + window_size / 2 - 1, y + window_size / 2 - 1)
538     half2 = fast_convolute(integral_image ,
539                           x - window_size / 2, y - window_size / 2,
540                           x - 1, y + window_size / 2 - 1)
541     return half1 - half2
542
543
544 def get_dy(integral_image , x, y, window_size):
545
546     half1 = fast_convolute(integral_image ,
547                           x - window_size / 2, y - window_size / 2 + 1,
548                           x + window_size / 2 - 1, y)
549     half2 = fast_convolute(integral_image ,
550                           x - window_size / 2, y + 1,
551                           x + window_size / 2 - 1, y + window_size / 2)
552     return half1 - half2
553
554
555 def is_corner(x, y, ratio_image , harris_window):
556
557     # Get the ratio of the pixel at (x, y).
558     ratio = ratio_image[y][x]
559
560     # For each pixel in the Harris window around the pixel (x,y).
561     for j in range(y - harris_window / 2,
562                  y + harris_window / 2 + 1):
563         for i in range(x - harris_window / 2,
564                       x + harris_window / 2 + 1):
565
566             # If the ratio of the pixel at (x, y) is less than one of its
567             # neighbors, eliminate it from the corners.
568             if ratio_image[j][i] > ratio:
569                 return False
570
571     return True
572
573
574 def get_shape_vectors(labels_image , components_corners):
575
576     # The unique labels.
577     labels = np.unique(labels_image)[1:]
578
579     # The array for the corners of each component.
580     components_centers = []
581
582     # The list of shape vectors of the components.
583     components_vectors = [[] for i in xrange(len(labels))]
584
585     # For each label:
586     for i, label in enumerate(labels):
587         # Find where the label is.
588         py, px = np.where(labels_image == label)
589
590         # Find the center of the component with that label.
591         center_x = (np.max(px) + np.min(px)) / 2
592         center_y = (np.max(py) + np.min(py)) / 2
593
594         # Add the center to the array of centers.
595         center = (center_x, center_y)
596         components_centers.append(center)
597
598         # The list of angles of the corners.
599         angles = []
600

```



```

601     # Compute the angles of the corners,
602     for corner in components_corners[i]:
603
604         angle = atan2(corner[1] - center[1], corner[0] - center[0])
605         angle = angle * 180 / pi
606         if angle < 0:
607             angle += 360
608
609         angles.append(angle)
610
611     # Sort the angles ascending.
612     angles.sort()
613
614     # Find the differences between the angles, which represent the arc
615     # length between the corners on the unit circle.
616     for j, angle in enumerate(angles):
617         diff = angle - angles[j - 1]
618         if diff < 0:
619             diff += 360
620         components_vectors[i].append(diff)
621
622     if len(components_vectors[i]) < CORNERS_NUM:
623         for j in xrange(CORNERS_NUM - len(components_vectors[i])):
624             components_vectors[i].append(0.0)
625
626     return components_centers, components_vectors
627
628
629 def recognize_components(training_vectors, components_vectors):
630
631     # The list of the best matches.
632     matches = []
633
634     # For each component:
635     for i, component_vector in enumerate(components_vectors):
636
637         # The index of the best match.
638         best_match = -1
639         # The minimum distance to the best match.
640         min_dist = sys.float_info.max
641
642         # For each element in the shape vector:
643         for j in xrange(len(component_vector)):
644
645             # Circularly rotate the vector around index j.
646             rotated_vector = component_vector[j:] + component_vector[:j]
647             #print rotated_vector
648
649             # For each training vector:
650             for k, training_vector in enumerate(training_vectors):
651
652                 #print training_vector
653                 # Calculate the euclidean distance.
654                 dist = np.linalg.norm((np.array(rotated_vector) -
655                                         np.array(training_vector)))
656
657                 # Keep the best match with its distance.
658                 if dist < min_dist:
659                     min_dist = dist
660                     best_match = k
661
662             matches.append(best_match)
663
664     return matches
665
666
667 if __name__ == "__main__":

```

```
668     main()
```

Listing 1: The entire Python source code