

ECE661: Homework 8

Fall 2016

Vinoth Venkatesan
venkat26@purdue.edu

November 7, 2016

Task 1.

Problem description: Implement the Local Binary Pattern (LBP) algorithm for texture characterization of images for the given classes of training images to obtain the texture histograms. Provide labels to histograms of pictures in the same class.

Solution:

The texture characterization of images can be done using various methods like GLCM (Gray Level Co-occurrence Matrix) and LBP (Local Binary Pattern). The implementation of LBP is explained below:

- **Finding coordinates and intensities of neighboring points:**

The neighbors of the pixel under construction is calculated using the change in distance along two orthogonal axis as:

$$(\Delta u, \Delta v) = \left(R \cos\left(\frac{2\pi p}{P}\right), R \sin\left(\frac{2\pi p}{P}\right) \right) \quad p = 0, 1, 2, \dots, 7$$

Here, the neighbors are found in a unit circle around the center pixel. The number of points is set by setting the value of $P = 8$. After finding the position of these pixels, we need to find the intensity of these point locations.

- If these points coincide with an existing pixel value, then the intensity of that pixel value is set to that point.
- If the point does not coincide with any pixel, we can use an interpolation technique. Here, a weighted mean of the intensity values of the four surrounding pixels was used to find the value of the current pixel as follows (here the weights are the inverse of the euclidean distances from the point):

$$I_i = \frac{I_{1i}w_1 + I_{2i}w_2 + I_{3i}w_3 + I_{4i}w_4}{w_1 + w_2 + w_3 + w_4}$$

- **Representing the intensities using a minimum value pattern**

In order to get the binary pattern which will act as a representation of these values, we use the following logic:

- If the intensity value is greater than or equal to the value at the center, we give it a label 1
- If the value is lesser, we give it a label 0.

This way, we get the binary pattern for the each pixel in the image. Now, in order to make the algorithm robust enough to reject in-plane rotations, we rotate the obtained binary pattern one bit at a time and keep doing it until we get the maximum number of zeros in the most significant bit (MSB) positions. This can done as follows:

- We rotate the pattern one bit at a time and find the unsigned integer value corresponding to that pattern.
- We use the pattern which has the least integer value and this pattern will represent the local texture of the pixel.

- **Encoding the binary pattern values:**

After we have acquired the binary pattern, we encode or represent these texture patterns based on the number of runs of 1's and 0's as follows:

- If there are exactly two runs, a run of 0's followed by 1's, represent the pattern by the number of 1's in the second run.
- If the pattern has all 0's, then represent it with 0.
- If the pattern has all 1's, represent it with the value $P(= 8)$.
- If the pattern has more than two runs, encode it with the value of P .

Once we have encoded the pixel, we update the histogram for the image which has 8-bins (equal to the value of P). This histogram then represents the image and will be used to compare with the test images to determine the label for the test image.

The LBP algorithm is implemented for all the training images in the four classes and each of these histograms are associated with a label depending on the class they belong to.

Task 2.

Problem description: Using the NN-Classifier, we find the label or the class that each of the test images belong to.

Solution:

The NN-Classifier uses the euclidean distance measure to classify the test images. The flow is explained as follows:

- For each of the test images, find the histogram texture for each of them using the LBP algorithm.
- Using Euclidean distance metric (2 - norm), we find the k -nearest neighbors by finding the metric value with each of the histograms of the training images. Here, the value of k was set to 5.
- We assign the corresponding label of the training image to the test image and hence we end up with k -labels. Out of these, we choose the label which occurs the most number of times and assign it to the test image.

Task 3.

Problem description: Based on the results obtained using the k -nearest neighbor algorithm, we build the confusion matrix to study the performance of the algorithm.

Solution:

The confusion matrix is built based on the classifications done for the test images. Hence, there are four rows and four columns (since we have four classes). The rows correspond to the actual labels and the columns correspond to the predicted labels. The confusion matrix obtained for the implementation is shown below:

	Car	Building	Mountain	Tree
Car	3	2	0	0
Building	0	4	1	0
Mountain	0	1	4	0
Tree	1	1	0	3

The accuracy of the algorithm was calculated to be 70%. The metric that was used is as follows:

$$Accuracy = \frac{Correct_Classifications}{Total_Number_of_Classifications}$$

Observations on performance of the ICP algorithm:

- The image recognition algorithm developed based on LBP and the NN-Classifier performed comparatively well given the accuracy of 70%.
- However, there were some wrong classifications and this could be due to a lot of reasons like:
 - The number of nearest neighbors which we considered were too small. This is a crucial one because it was observed that in a lot of cases, there were equal number of wrong neighbors and correct neighbors it is difficult to automate the decision making process in these cases. Hence, increasing the number k might result in a better performance.
 - The resolution of the pictures in both the training and test images were small and this could have resulted in wrong classifications.
 - The encoding algorithm which we used seems to be biased towards the last condition (where there are more than two runs) which was apparent from the histograms. If we use a better encoding technique, then we might get better results.

Results:




Figure 1: LBP histogram for a car image

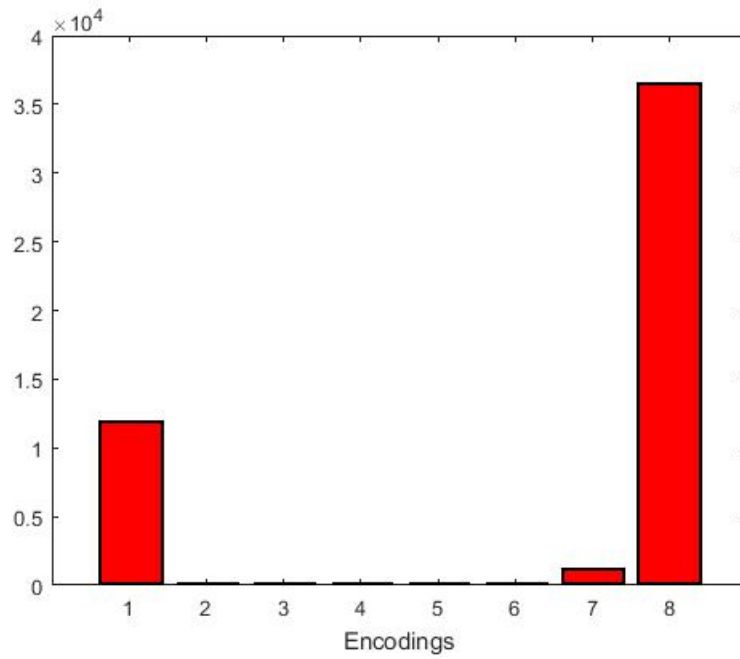


Figure 2: LBP histogram for a building image

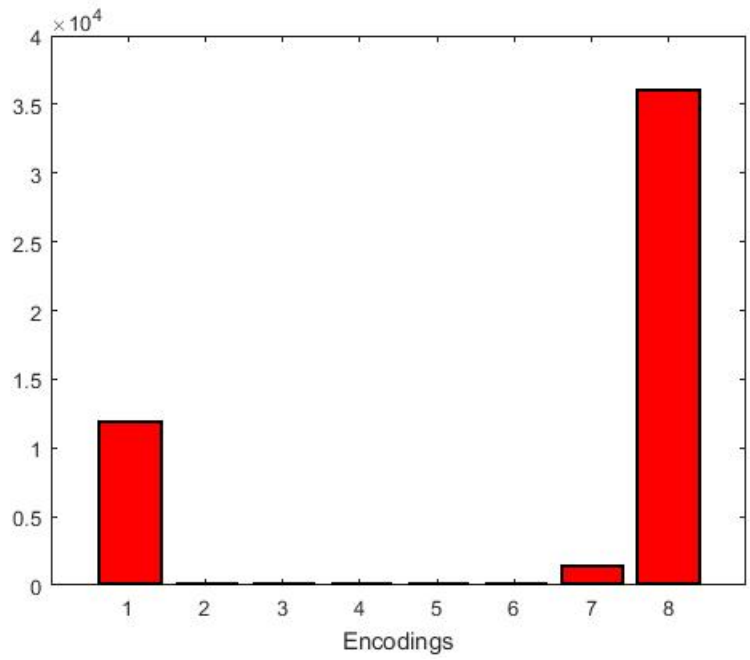


Figure 3: LBP histogram for a mountain image




Figure 4: LBP histogram for a tree image

Source Code(Main):

```
1 clear all
2 close all
3 clc
4
5 for i = 1:20
6
7     % Load i-th image in the training set
8     str = sprintf('01_car (%d).jpg', i);
9     A = imread(str);
10    % Convert to grayscale
11    A = rgb2gray(A);
12    % Call to LBP_texture to get the texture histogram
13    histogram_car(i,:) = LBP_texture(A);
14
15    str = sprintf('01_building (%d).jpg', i);
16    A = imread(str);
17    A = rgb2gray(A);
18    histogram_building(i,:) = LBP_texture(A);
19
20    str = sprintf('01_mountain (%d).jpg', i);
21    A = imread(str);
22    A = rgb2gray(A);
23    histogram_mountain(i,:) = LBP_texture(A);
24
25    str = sprintf('01_tree (%d).jpg', i);
26    A = imread(str);
27    A = rgb2gray(A);
28    histogram_tree(i,:) = LBP_texture(A);
29
30 end
31
32 % Saving the training results
33 save('training_results_LBP');
```

LBP_texture

```
1
2 function histogram = LBP_texture(A)
3
4 R = 1; % Radius for examination
5 P = 8; % Number of sampled
6 histogram = zeros(1,8);
7
8 for i = 2:size(A,1)-1
9     for j = 2:size(A,2) - 1
10
```

```

11 % Pixel coordinates of the neighbours
12 neighbours = [i + R*(cos((2*pi*(0:7))/P)); j + R*(sin((2*pi
13 * (0:7))/P))];
14 intensity = zeros(1,8);
15
16 % Finding the intensity of the neighbouring pixels
17 for k = 1:size(neighbours,2)
18     if (mod(k,2))
19         % When the point coincides with a pixel coordinate
20         if(A(round(neighbours(1,k)),round(neighbours(2,k)))
21             >= A(i,j))
22             intensity(k) = 1;
23         else
24             intensity(k) = 0;
25         end
26     else
27         % Performing weighted mean to find the intensity
28         if (weighted_mean(neighbours(2,k),neighbours(1,k),A)
29             >= A(i,j))
30             intensity(k) = 1;
31         else
32             intensity(k) = 0;
33         end
34     end
35 end
36
37 % Initializations
38 final_num = intensity;
39 bin_num = intensity;
40 final_decimal = num2str(intensity);
41 final_decimal(isspace(final_decimal)) = '';
42 final_decimal = bin2dec(final_decimal);
43
44 % Loop to find the minimum value pattern
45 for k = 1:8
46     % Rotating one bit at a time
47     bin_num = circshift(bin_num,1,2);
48     str_x = num2str(bin_num);
49     str_x(isspace(str_x)) = '';
50     bin_num = bin2dec(str_x);
51
52     % Finding the minimum value
53     if bin_num < final_decimal
54         final_decimal = bin_num;
55         final_num = de2bi(bin_num,8,'left-msb');
56     end
57 end

```



```

55         bin_num = de2bi(bin_num,8,'left-msb');
56     end
57
58     % Initialize runs to zero
59     runs = 0;
60
61     % Loop to find the number of runs
62     for k = 1:7
63         if final_num(k + 1) ~= final_num(k)
64             runs = runs + 1;
65         end
66     end
67
68     % Updating the histogram vector based on the number of runs
69     if ((runs == 0) && (final_num(1) == 1)) || runs > 1
70         histogram(8) = histogram(8) + 1;
71     elseif (runs == 0) && (final_num(1) == 0)
72         histogram(1) = histogram(1) + 1;
73     else
74         temp = nnz(final_num == 1);
75         histogram(temp) = histogram(temp) + 1;
76     end
77
78 end
79 end
80
81
82 end

```

main_testing.m

```

1  clear all
2  close all
3  clc
4
5  % Load the training results
6  load training_results_LBP
7  histogram_all = cat(1,histogram_car,histogram_building,
8                     histogram_mountain,histogram_tree);
9
10
11  for i = 1:5
12
13      % Loading the i-th test image
14      str = sprintf('car_%d.jpg', i);
15      A = imread(str);
16      % Convert to grayscale
17      A = rgb2gray(A);
18      % Finding the histogram of the test image

```

```

17 histogram_car_test(i,:) = LBP_texture(A);
18 % Classification based on the NN-Classifier
19 label_test_car(i) = NN_classifier(histogram_car_test(i,:),
    histogram_all);
20
21 str = sprintf('building_%d.jpg', i);
22 A = imread(str);
23 A = rgb2gray(A);
24 histogram_building_test(i,:) = LBP_texture(A);
25 label_test_building(i) = NN_classifier(histogram_building_test(i
    ,:), histogram_all);
26
27 str = sprintf('mountain_%d.jpg', i);
28 A = imread(str);
29 A = rgb2gray(A);
30 histogram_mountain_test(i,:) = LBP_texture(A);
31 label_test_mountain(i) = NN_classifier(histogram_mountain_test(i
    ,:), histogram_all);
32
33 str = sprintf('tree_%d.jpg', i);
34 A = imread(str);
35 A = rgb2gray(A);
36 histogram_tree_test(i,:) = LBP_texture(A);
37 label_test_tree(i) = NN_classifier(histogram_tree_test(i,:),
    histogram_all);
38
39 end
40
41 % Confusion matrix development
42 confusion_matrix = [nnz(label_test_car == 1)    nnz(label_test_car
    == 2)    nnz(label_test_car == 3)    nnz(label_test_car == 4)
43    nnz(label_test_building == 1)    nnz(label_test_building == 2)
    nnz(label_test_building == 3)    nnz(label_test_building
    == 4)
44    nnz(label_test_mountain == 1)    nnz(label_test_mountain == 2)
    nnz(label_test_mountain == 3)    nnz(label_test_mountain
    == 4)
45    nnz(label_test_tree == 1)    nnz(label_test_tree == 2)    nnz(
    label_test_tree == 3)    nnz(label_test_tree == 4)];
46
47 % Accuracy = correct_classifications/total_number_of_classifications
48 Accuracy = trace(confusion_matrix)/sum(sum(confusion_matrix));

```

NN_classifier.m

```

1 function label = NN_classifier(histogram_test, histogram_all)
2
3 % Initialize the distance matrix

```

```

4 euclidean_dist = zeros(1,size(histogram_all,1));
5
6 % Finding the euclidean distance
7 for i = 1:size(histogram_all,1)
8     euclidean_dist(i) = norm(histogram_test - histogram_all(i,:));
9 end
10
11 % Finding the k-nearest neighbours (for 5 neighbours)
12 [minimum, indices] = sort(euclidean_dist);
13 labels = zeros(1,5);
14
15 for i = 1:5
16     if indices(i) <= 20
17         labels(i) = 1;
18     elseif indices(i) <= 40
19         labels(i) = 2;
20     elseif indices(i) <= 60
21         labels(i) = 3;
22     else
23         labels(i) = 4;
24     end
25 end
26
27 % Finding the label for the test image
28 temp = [nnz(labels == 1) nnz(labels == 2) nnz(labels == 3) nnz(
29     labels == 4)];
30 % label that appears the most
31 [maximum, label] = max(temp);
32 end

```

weighted_mean.m

```

1 function X = weighted_mean(x,y,B)
2
3     % Finding the surrounding pixels
4     x1 = floor(x);
5     x2 = ceil(x);
6     y1 = floor(y);
7     y2 = ceil(y);
8
9     % Finding the weights (euclidean distance)
10    dist1 = norm([x - x1, y - y1]);
11    dist2 = norm([x - x2, y - y1]);
12    dist3 = norm([x - x1, y - y2]);
13    dist4 = norm([x - x2, y - y2]);
14
15    y_act1 = size(B,1) - y1 + 1;

```

```
16     y_act2 = size(B,1) - y2 + 1;
17     x_act1 = x1;
18     x_act2 = x2;
19
20     % Calculating the intensity value using weighted mean
21     X = (B(y_act1,x_act1,:)./dist1 + B(y_act1,x_act2,:)./dist2 + B(
22         y_act2,x_act1,:)./dist3 + B(y_act2,x_act2,:)./dist4)...
23         ./((1/dist1 + 1/dist2 + 1/dist3 + 1/dist4));
24 end
```
