

CSE152A – Computer Vision – Assignment 2 (SP16)

Instructor: Ben Ochoa

Maximum Points : 65

Deadline : 11:59 p.m., Wednesday, 04-May-2016

Instructions:

- This assignment must be solved, and written up individually
- There is no physical hand-in for this assignment.
- Coding for this assignment must be done in MATLAB. In general, code does not have to be efficient. Focus on clarity, correctness and functionality here, we can worry about speed in another course.
- All code developed for this assignment must be included in the appendix of the report. If the appendix does not contain your code, points may be deducted.
- Submit your assignment electronically by email to Akshat Dave [akdave@ucsd.edu] and Mihir Patankar [mpatanka@eng.ucsd.edu] with the subject line *CSE152-Assignment-2*. The email should contain one attached file named [`CSE_152_HW2_<student-pid>.zip`]. This zip file must contain the following two artifacts:
 1. A pdf file named [`CSE_152_HW2_<student-pid>.pdf`] containing your writeup. Please clearly state the author's full name and student identity in the report. The report must contain all the result images. Images not embedded in the report will not be considered for evaluation. In the case of figures containing text and labels, it is your responsibility to ensure that the text is readable; points may be deducted if these are not readable.
 2. A folder named [`CSE_152_HW2_<student-pid>.code`] containing all your matlab code files

Introduction

In this assignment, we will look at binary image processing, connected component analysis and image filtering.

1 Binarization [10 points]

Write a MATLAB function to implement Otsu's method described in class (CSE152-notes-link). This algorithm should automatically determine an intensity level to threshold an image to segment out the foreground from the background. The output of your function should be a binary image that is 0 for all background pixels and 1 for all foreground pixels. Apply this function to the image `can_pix.png` and turn in the output image in your report.

Notes:

- Load in an image and convert it to grayscale using `rgb2gray.m`.
- Normalize the image and convert it to double precision using `im2double.m`
- You can use the MATLAB function `hist.m` to create a histogram of pixel intensities as a first step.

2 Connected Component Analysis [40 points]

In this section, we will look at connect components in images and perform some analysis on these components. Using certain properties of the regions (discussed in lecture), we will proceed to rectify the foreground components.

2.1 Connected Regions [15 points]

- Write **MATLAB** code to implement the connected component labeling algorithm discussed in class, assuming 8-connectedness. Your function should take as input a binary image (computed using your algorithm from question 1) and output a 2D matrix of the same size where each connected region is marked with a distinct positive number (e.g. 1, 2, 3). On the image `can_pix.png`, display an image of detected connected components where connected region is mapped to a distinct color (using the function `imagesc` in **MATLAB**). You may need to modify the **MATLAB** recursion limits using `set(0,'RecursionLimit',1000)`
- How many components do you think are in the image `coins_pix.jpg`? What does your connected component algorithm give as output for this image? Include your output on this image in your report. it may help to reduce the size of the image before running the connected components algorithm but after converting the image to a binary image `imresize(binary, 0.25, 'bilinear')`

2.2 Take your own images [5 points]

For this part, you will be using images you take with your own camera. Choose three objects of different shapes, preferably with non-shiny surfaces. Possibilities include paper or cardboard cut-outs, bottle tops, pencils, Lego pieces, etc. Take 3 pictures of these objects individually with a solid background. The object should be clearly distinguishable from the background (e.g. bright object & dark background). Include these three images in your report as in Figure 1. In addition, show similar output images as in Problem 2.1 for each of your new images (there is only 1 connected component in this case).

2.3 Image moments and rectification [10 points]

Write three functions which compute the moments, central moments, and normalized moments of a marked region (CSE152-notes). Each function should take as input a 2D matrix (output of part 2.2) and 3 numbers j , k , and d . The output should be the (j, k) moment $M_{j,k}$, central moment $\mu_{j,k}$, normalized moment $m_{j,k}$ of the region marked with positive number d in the input matrix.

Using these functions, on each of the three images, draw the centroid of each object. Also compute the eigenvectors of the centralized second moment matrix (CSE152-notes) and draw the two eigenvectors on the centroid. This should indicate the orientation of each object. See Figure 2 for the results on the example images. Turn in the outputs for your own images.



Figure 1: Sample images

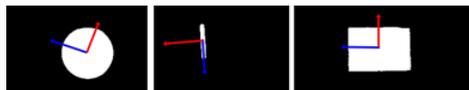


Figure 2: Sample images with principal directions

2.4 Image alignment [10 points]

The orientation computed from Problem 2.3 can be used to roughly align the orientation of the region (i.e. in-plane rotation). Write a function to rotate the region around its centroid so that the eigenvector corresponding to the largest eigenvalue (i.e. the blue vector in Figure 3) will be horizontal (aligned with $[1, 0]^T$). This might look like like Figure 3.

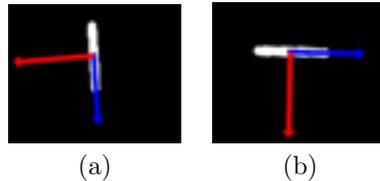


Figure 3: (a) Original binarized image with principal directions (b) Aligned sample image with major eigenvector aligned along $[1, 0]^T$

Your function should take as input a 2D matrix (output of Problem 2.1) and output a matrix of the same size in which all marked regions are aligned. Turn in the aligned outputs for your images.

Note: After finding the rotation matrix R to rotate the largest eigenvector to $[1, 0]^T$, we rotate all points (x, y) belonging to that region using the following transformation

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = R \begin{bmatrix} x - \hat{x} \\ y - \hat{y} \end{bmatrix} + \begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix}$$

where $[\hat{x}, \hat{y}]^T$ are the centroid coordinates of the region. For simplicity, just ignore the cases when part of the aligned region falls outside of the image border or is overlapped with other regions. You can avoid these issues when capturing your images (e.g. put your objects a bit far apart). Finally, note that the rotation matrix can be created trivially from the eigenvectors.

3 Filtering [15 points]

In this problem we will play with convolution filters. Filters, when convolved with an image, will respond strongest on locations of an image that look like the filter. This allows us to use filters as object templates in order to identify specific objects within an image. In the case of this assignment, we will be finding cars within an image by convolving a car template onto that image. Although this is not a very good way to do object detection, this problem will show you some of the steps necessary to create a good object detector. The goal of this problem will be to teach some pre-processing steps to make vision algorithms be successful and some strengths and weaknesses of filters. Each problem will ask you to analyze and explain your results. If you do not provide an explanation of why or why not something happened, then you will not get full credit.

3.1 Warmup - Mickey Detection [5 points]

First you will convolve a filter to a synthetic image. The filter or template is `filter.jpg` and the synthetic image is `toy.png`. These files are available on the course webpage. You will first want to modify the filter image and original slightly. We will do so by subtracting the mean of the image intensities from the image i.e. $I_t \leftarrow I - \text{mean}(\text{vec}(I))$, where I_t is the transformed image and I is the original image. To convolve the filter image with the toy example, in MATLAB you will want to use `imfilter`. The output of the convolution will create an intensity image. Provide this image in the report. In the original image (not the image with its mean subtracted), draw a bounding box of the same size as the filter image around the top 3 intensity value locations in the convolved image.

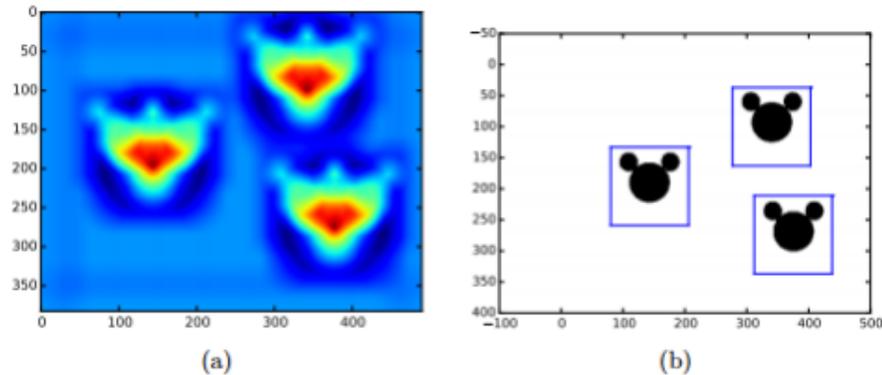


Figure 4: Example outputs for the synthetic example. (a) Heat map. (b) Bounding boxes

The outputs should look like Figure 4. Describe how well you think this technique will work on more realistic images? Do you foresee any problems for this algorithm on more realistic images?

3.2 Detection quality [4 points]

We have now created an algorithm that produces a bounding box around a detected object. However we have no way to know if the bounding box is good or bad. In the example images shown above, the bounding boxes look reasonable, but not perfect. Given a ground truth bounding box (g) and a predicted bounding box (p), a commonly used measurement for bounding box quality is:

$$r = \frac{p \cap g}{p \cup g}$$

More intuitively, this is the number of overlapping pixels between the bounding boxes divided by the total number of unique pixels of the two bounding boxes combined. Assuming that all bounding boxes will be squares (and not diamonds), implement this error function. Report the detection ratio r values (as a percentage) for the bounding box specified in the “Mickey Detection” problem. In general, if the overlap is 50% or more, you may consider that the detection did a good job.

3.3 Car Detection [6 points]

Now that you have created an algorithm for matching templates and a function to determine the quality of the match, it is time to try some more realistic images. The file, `cartemplate.jpg`, will be the filter to convolve on each of the 5 other car images (`car1.jpg`, `car2.jpg`, `car3.jpg`). Each image will have an associated text files that contains two x, y coordinates (one pair per line). These coordinates will be the ground truth bounding box for each image. For each car image, provide the following:

- (i) A heat map image
- (ii) A bounding box drawn on the original image.
- (iii) The bounding box overlap percent $r \times 100\%$.

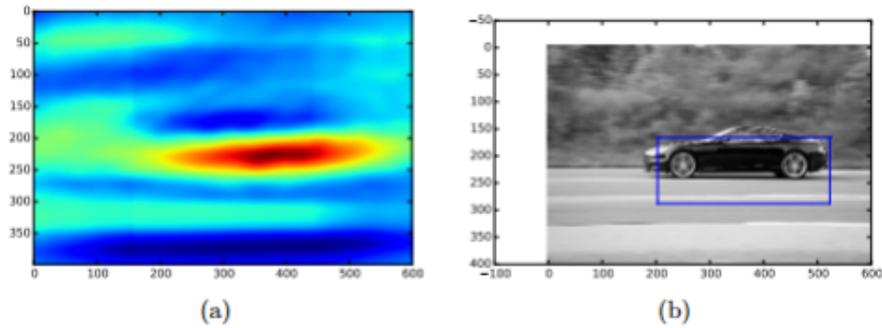


Figure 5: Example outputs for the car example. (a) Heat map. (b) Bounding boxes

Here are some useful hints to increase the overlap percentage:

- Rescaling the car template to various sizes (for `car1.jpg`)
- Horizontally flipping the car template (for `car2.jpg`)
- A combination of the first two hints (for `car3.jpg`)
- Gaussian blurring might be useful in all cases

An example output is show for `car1.jpg` in Figure 5. It may not be possible to achieve 50% overlap on all the images. Your analysis of the images will be worth as much as achieving a high overlap percentage.