

<http://crcv.ucf.edu/people/faculty/Bagci/>

[PROGRAMMING ASSIGNMENT] (1)

COMPUTER VISION

DR. ULAS BAGCI • (FALL) 2015 • UNIVERSITY OF CENTRAL FLORIDA (UCF)

Coding Standard and General Requirements

Code for all programming assignments should be **well documented**. A working program with no comments will receive **only partial credit**. Documentation entails writing a description of each function/method, class/structure, as well as comments throughout the code to explain the program flow. Programming language for the assignment is **Python**. You can use standard python built-in IDLE, or CANOPY for the working environment. Other commonly used IDLEs are the following: PyCharm Community Edition, PyScripter, CodeSculptor, Eric Python, Eclipse plus PyDev.

Following libraries will be used extensively throughout the course:

- PIL (The Python Imaging Library), Matplotlib, NumPy, SciPy, LibSVM, OpenCV, VLFeat, python-graph.

If you use CANOPY, make sure that you use version 2.7, which already includes many libraries. If you are asked to implement “Gaussian Filtering”, you are not allowed to use a Gaussian function from a known library, you need to implement it from scratch.

Submit by **10th of September 2015**, 11.59pm.

Question 1: Canny Edge Detection Implementation

[4 pts] In 1986, John Canny defined a set of goals for an edge detector and described an optimal method for achieving them. Canny specified three issues that an edge detector must address:

- **Error rate:** Desired edge detection filter should find all the edges, there should not be any missing edges, and it should respond only to edge regions.
- **Localization:** Distance between detected edges and actual edges should be as small as possible.
- **Response:** The edge detector should not identify multiple edge pixels where only a single edge exists.

Remember from the lecture that in Canny edge detection, we will first smooth the images, then compute gradients, magnitude, and orientation of the gradient. This procedure is followed by non-max suppression, and finally hysteresis thresholding is applied to finalize the steps. Briefly, follow the steps below for practical implementation of *Canny Edge detector*:

1. Read a gray scale image you can find from [Berkeley Segmentation Dataset, Training images](#), store it as a matrix named I .
2. Create a one-dimensional Gaussian mask G to convolve with I . The standard deviation(s) of this Gaussian is a parameter to the edge detector (call it $\sigma > 0$).


3. Create a one-dimensional mask for the first derivative of the Gaussian in the x and y directions; call these G_x and G_y . The same $\sigma > 0$ value is used as in step 2.
4. Convolve the image I with G along the rows to give the x component image (I_x), and down the columns to give the y component image (I_y).
5. Convolve I_x with G_x to give I'_x , the x component of I convolved with the derivative of the Gaussian, and convolve I_y with G_y to give I'_y , y component of I convolved with the derivative of the Gaussian.
6. Compute the magnitude of the edge response by combining the x and y components. The magnitude of the result can be computed at each pixel (x, y) as: $M(x, y) = \sqrt{I'_x(x, y)^2 + I'_y(x, y)^2}$.
7. Implement **non-maximum suppression** algorithm that we discussed in the lecture. Pixels that are not local maxima should be removed with this method. In other words, not all the pixels indicating strong magnitude are edges in fact. We need to remove false-positive edge locations from the image.
8. Apply **Hysteresis thresholding** to obtain final edge-map.

Definition: Non-maximal suppression means that the center pixel, the one under consideration, must have a larger gradient magnitude than its neighbors in the gradient direction. That is: from the center pixel, travel in **the direction of the gradient** until another pixel is encountered; this is the first neighbor. Now, again starting at the center pixel, travel in the direction opposite to that of the gradient until another pixel is encountered; this is the second neighbor. Moving from one of these to the other passes through the edge pixel in a direction that crosses the edge, so the gradient magnitude should be largest at the edge pixel.

Algorithmically, for each pixel \mathbf{p} (at location x and y), you need to test whether a value $M(\mathbf{p})$ is maximal in the direction $\theta(\mathbf{p})$. For instance, if $\theta(\mathbf{p}) = \pi/2$, i.e., the gradient direction at $\mathbf{p} = (x, y)$ is downward, then $M(x, y)$ is compared against $M(x, y - 1)$ and $M(x, y + 1)$, the values above and below of \mathbf{p} . If $M(\mathbf{p})$ is not larger than the values at both of those adjacent pixels, then $M(\mathbf{p})$ becomes 0. For estimation of the gradient orientation, $\theta(\mathbf{p})$, you can simply use $\text{atan2}(I'_y, I'_x)$.

Hint: It is assumed that the gradient changes continuously as a function of position, and that the gradient at the pixel coordinates are simply sampled from the continuous case. If it is further assumed that the change in the gradient between any two pixels is a linear function, then the gradient at any point between the pixels can be approximated by a linear interpolation.

For a sample output, please refer to figure below: chessboard image's X component of the convolution with a Gaussian (a), Y component of the convolution with a Gaussian (b), X component of the image convolved with the derivative of a Gaussian (c), Y component of the image convolved with the derivative of a Gaussian (d), resulting magnitude image (e), and canny-edge image after non-maximum suppression (f) are shown.

**Your tasks:**

- Choose three example gray-scale images from Berkeley Segmentation Dataset (Training Images) [CLICK HERE](#). When executed, your algorithm should plot intermediate and final results of Canny Edge Detection process as similar to the figure illustrated above.
- Please show the effect of σ in edge detection by choosing three different σ values when smoothing. Note that you need to indicate which σ works best as a comment in your assignment.

Question 2: Quantitative Evaluation of Edge Detector

[1 pt] Download the three input images (and corresponding edge maps of those three images) from the webcourses. After using your own implementation of Canny edge detection output on input images, you will have three output images. Use the following metrics to evaluate your edge detector with respect to the edge maps provided to you along with input images.

Your task is to report quantitative evaluation of your edge detection results with the following functions:

- **Sensitivity** (or true positive rate TPR) as $TP/(TP + FN)$
- **Specificity** (or true negative rate TNR) as $TN/(TN + FP)$
- **Precision** (or positive predictive value PPV) as $TP/(TP + FP)$
- **Negative Predictive Value** as $TN/(TN + FN)$
- **Fall-out** (or false positive rate FPR) as $FP/(FP + TN)$
- **False Negative Rate FNR** as $FN/(FN + TP)$

- **False Discovery Rate FDR** as $FP/(FP + TP)$
- **Accuracy** as $(TP + TN)/(TP + FN + TN + FP)$
- **F-score** as $2TP/(2TP + FP + FN)$
- **Matthew's Correlation Coefficient** as $\frac{TP*TN-FP*FN}{\sqrt{((TP+FP)*(TP+FN)*(TN+FP)*(TN+FN))}}$

[1 pt] Select the first input image, and add **Gaussian** and **Salt-Pepper** noise to the input image. Now, **your task** is to repeat the previous step (find Canny edges of the noisy images), and compute the aforementioned metrics with respect to the ground truth image. For Gaussian noise parameters, the mean should be set to 0 and standard deviation be 10. For Salt-pepper noise, the amount of standard deviation is interpreted as a percentage, and it should be set to 25%.

Question 3: Residual Images and Smoothing

[1 pt] Use a 3 x 3 box filter recursively (up to 30 iterations) for generating residual images with respect to smoothing. Compare with residual images when smoothing with the Gauss filter of size $(2k+1) \times (2k+1)$ for $k = 1, \dots, 15$. Discuss the general relationship between recursively repeated box filters and a Gauss filter of the corresponding radius. Actually, what is the corresponding radius? Write your answer as a comments, and show necessary plots for this question.

Question 4: Can we use Entropy for thresholding?

[3 pts] This question examines your knowledge on probability density function (pdf) construction from gray-scale histogram, and the use of Entropy information for mapping image intensity values into two classes: white (foreground) and black (background).

Definition 1: Entropy (i.e., H) is a measure of disorder (uncertainty), and it is zero for a perfectly ordered system. The entropy statistic is high if a variable is well distributed over the available range, and low if it is well ordered and narrowly distributed. (Entropy formulation is given in Lecture 3 notes).

Definition 2: The **Histogram** h of a digital image I is a plot or graph of the frequency of occurrence of each gray level in I . Hence, h is a one-dimensional function with domain $0, \dots, L - 1$ and possible range extending from 0 to the number of pixels in the image, N .

Histogram and *pdf* (probability density function) have the following direct relationship. Assume that the image I has a grayscale resolution of L gray levels. The number of pixels with gray level i is written as n_i , so the total number of pixels in the image is $N = n_0 + n_1 + \dots + n_{L-1}$. Thus, the probability of a pixel having gray level i is:

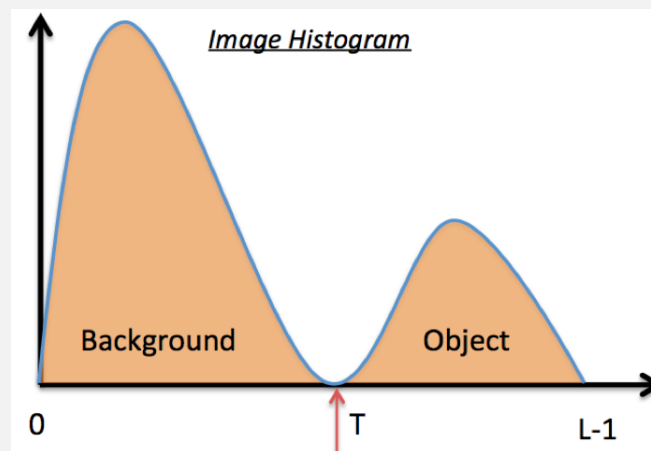
$$p_i = \frac{n_i}{N},$$

where $p_i \geq 0$ and $\sum_{i=1}^L p_i = 1$. Thus, normalized histogram is simply *pdf*.

Definition 3: Thresholding is an extreme form of gray level quantization. Suppose that a gray level image I can take L possible gray levels (i.e., for 8-bit, $L=256$). The process of thresholding is a simple comparison: each pixel value in I is compared to T , where $T \in 0, \dots, L-1$. Based on this comparison, a binary decision is made that defines the value of the corresponding pixel (at location x,y) in an output binary image:

$$O(x, y) = \begin{cases} 1 & \text{if } I(x, y) \geq T \\ 0 & \text{otherwise.} \end{cases}$$

The threshold T is of critical importance, since it controls the particular abstraction of information that is obtained. Indeed, different thresholds can produce different valuable abstractions of the image. Image thresholding can often produce a binary image result that is quite useful for simplified processing, interpretation, or display. However, some gray level images do not lead to any interesting binary result regardless of the chosen threshold T . Several questions arise: given a gray level image, how can one decide whether binarization of the image by gray level thresholding will produce a useful result? Can this be decided automatically by a computer algorithm? Assuming that thresholding is likely to be successful, how does one decide on a threshold level T ? These are apparently simple questions pertaining to a very simple operation. However, the answers to these questions turn out to be quite difficult to answer in the general case. In all cases, however, the basic tool for understanding the process of image thresholding is the image histogram (see figure below).



The concept of entropy thresholding is to threshold at an intensity for which the sum of the entropies of the two intensity probability distributions thereby separated is maximized. The reason for this is to obtain the greatest reduction in entropy—i.e., the greatest increase in order—by applying the threshold: in other words, the most appropriate threshold level is the one that imposes the greatest order on the system, and thus leads to the most meaningful result. To proceed, the intensity probability distribution is again divided into two classes (i.e., A and B)—those with gray levels above the threshold value T and those with gray levels above T . This leads to two probability distributions A and B:

$$A : \frac{p_0}{P_T}, \frac{p_1}{P_T}, \dots, \frac{p_T}{P_T},$$

$$B : \frac{p_{T+1}}{1 - P_T}, \frac{p_{T+2}}{1 - P_T}, \dots, \frac{p_{L-1}}{1 - P_T}.$$

where $P_T = \sum_{i=0}^T p_i$.

Your tasks are:

- Choose three example gray-scale images from Berkeley Segmentation Dataset (Training Images) [CLICK HERE](#).
- For each possible $T \in 0, \dots, L-1$, compute summation of entropy A and entropy B , called *total entropy* $= H(A) + H(B)$. Note that A and B correspond to background and foreground of the image.
- Find the value of T corresponding to maximum total entropy $H(T) = H(A) + H(B)$.
- Plot each of three gray scale images (I) and corresponding binary images (O) on the screen (binary images are obtained through thresholding gray-scale images at threshold level T).