http://cs.ucf.edu/~bagci/

[Programming Assignment] (1)
Robot Vision

Dr. Ulas Bagci • (Spring) 2018 • 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 other IDLEs such as CANOPY, PyCharm Community Edition, PyScripter, CodeSculptor, Eric Python, Eclipse plus PyDev, etc.

Following libraries will be used extensively throughout the course:

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

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.

Along with your well commented code, please submit a "brief" report as doc or pdf format.

Submit by **1st of February 2018**, 11.59pm.

Question 1: Box Filtering [1 pt]

Implement two box filters (one with 3 by 3, the other one is 5 by 5 kernel size), and apply them to the given images (i.e., image1.png and image2.png) separately. Show the resulting images, and explain the resulting images, and their differences. (For convolution operation, you can use built-in function. Do not use built-in function for box filtering.)

Question 2: Median Filtering [1 pt]

Implement three median filters (3 by 3, 5 by 5, and 7 by 7 kernel size), and apply these filters to image1.png and image2.png separately. Discuss the resulting differences for each kernel, explain where median filters are most effective. (For convolution operation, you can use built-in function. Do not use built-in function for Median filtering.)

Question 3: Gaussian Filtering [1 pt]

Implement a two-dimensional Gaussian kernel with a variation (sigma) equal to 3, 5, and 10. Apply these three Gaussian kernels to image1.png and image2.png separately, discuss the differences of Gaussian operations with different sigmas. Also, compare your results with question 1 and question 2: what are the differences between these

three filters, what do you observe? Which filtering is most effective in which images ? Why ? (For convolution operation, you can use built-in function. Do not use built-in function for Gaussian filtering.)

Question 4: Gradient Operations [1 pt]

Write a derivative operations (forward, backward, and central difference) that can be applied to any given image (f) and produces two images: gradient x, and gradient y images (i.e., f_x, f_y). Also, calculate gradient magnitude as $\sqrt{(f_x^2 + f_y^2)}$. The image that you should use for this problem is called image4.png. (Do not use built-in gradient functions, create difference operator yourself as stated clearly in the question.)

Question 5: Sobel Filtering [1 pt]

Implement Sobel filtering with 3 by 3 kernel size. Apply it to image1.png and image2.png. Discuss the resulting images. (For convolution operation, you can use built-in function. Do not use built-in function for Sobel filtering.)

Question 6: Fast Gaussian Filtering [1 pt]

Implement question 3, but this time, use only 1D Gaussian to do filtering in 2D. The trick is the following: since Gaussians are separable, you can use 1D Gaussian to filter the image in one direction first, and then the other direction can be filtered. Compare the efficiency of both methods (question 3 and question 6). Use image1.png and image2.png for smoothing. (For convolution operation, you can use built-in function. Do not use built-in function for fast Gaussian filtering.)

Question 7: Histogram [1 pt]

Implement histogram function from scratch, and show the resulting bar-graph (histogram). Use 256, 128, and 64 bins to visualize histograms. Comment on the resulting differences with respect to bins. Use image4.png to conduct this experiment. (Do not use built-in histogram function, create it yourself.)

Question 8: Thresholding [1 pt]

Implement a simple thresholding function to create a binary image. Assume that you want to binarize an image by selecting a threshold value T such that all the intensity values in the image below T will be assigned 0, and all the intensity values in the image that are equal and greater than T will be assigned 1. Resulting image is called binary image (as it only includes 0 and 1), and while 0 means background, 1 indicates a foreground. Use image4.png for this operation, and ask users in your program to enter a value between 0 to 256 as a threshold value. After user entered a value for T, show the resulting image as an output. (Do not use built-in histogram function, create it yourself. See definition 3 below in question 9 for a simple formulation of thresholding operation.)

Question 9: Can we use Entropy for thresholding? [2 pt]

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 for a given N samples can be formulated as follows: $H = \sum_{i=1}^{N} p_i.log(p_i)$, where p_i means probability of the sample i.).

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 $0, ..., L - 1$ 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 $N = n_0 + n_1 + ... + n_{L-1}$. 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 $T \in 0, ..., L - 1$. 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).
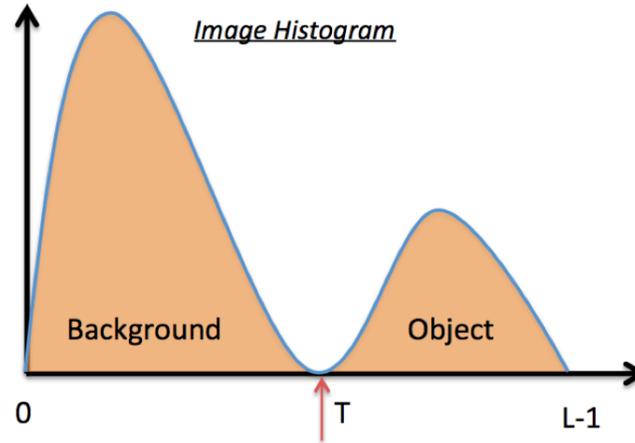
The concept of entropy thresholding is to threshold at an intensity for which the sum of the entropies of the two intensity probability distributions thereby separated is maximized. The reason for this is to obtain the greatest reduction in entropy–i.e., the greatest increase in order–by applying the 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:

$$A : \frac{p_0}{P_T}, \frac{p_1}{P_T}, ..., \frac{p_T}{P_T},$$
$$B : \frac{p_{T+1}}{1 - P_T}, \frac{p_{T+2}}{1 - P_T}, ..., \frac{p_{L-1}}{1 - P_T}.$$

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

Your tasks are:

- Use the attached image for conducting this experiment.

- For each possible $T \in 0, ..., L-1$, 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) .