

Intro to Image Understanding (CSC420)

Assignment 1

Posted: Jan 18, 2018 Submission Deadline : Jan 25, 11.59pm, 2018

Instructions for submission: Please write a document (either pdf, doc, etc) with your solutions (include pictures where needed). Include your code inside the document. Please submit through MarkUs. You are expected to work on the assignment **individually**.

Max points: 15, max extra credit points: 3

1. (a) **[2 points]** Write your own code for computing correlation of the 2D (grayscale) image and a 2D filter. Make the output matrix be the same size as the input image. Be careful to correctly deal with the border of the image – the easiest way to do this is to “zero-pad” the image prior to correlation.
(b) **[1 point]** Extend this code to handle RGB images and 3D filters (having the third dimension equal to 3).
2. (a) **[1 point]** Given a $n \times n$ image, I , and $m \times m$ filter, h , what is the computational cost of computing $h * I$ (convolution)? What is the computational cost if h is a separable filter?
(b) **[1 point]** Write your own function that creates an anisotropic Gaussian filter with σ_x and σ_y as an input parameter.
(c) **[1 point]** Convolve the attached `image.png` with a (2D) Gaussian filter with $\sigma_x = 1, \sigma_y = 10$ and visualize the result (display the result of the convolution). You can use built-in functions for convolution.
(d) **[1 point]** Is a horizontal derivative, $\frac{\partial G(x,y)}{\partial x}$, of a Gaussian filter G a separable filter? Explain your answer.
(e) **[1 point]** Given a filter F , how can you (automatically) check whether it's separable?
3. (a) **[1 point]** Compute magnitude of gradients for the attached images `waldo.png` and `template.png`.
(b) **[1 point]** Write a function that localizes the template (`template.png`) in the image `waldo.png` based on the magnitude of gradients. You can help yourself with the function available under Lecture 2 on class webpage (be careful, that function uses intensity values).

4. (a) [1 point] How does the Canny edge detector work? Please explain all the steps, and provide brief motivation for each step.
- (b) [1 point] Run the Canny edge detector on `court.jpg`. Play with the parameters so that you get rid of low-contrast edges. In Matlab you can help yourself with function `edge`.
5. **Extra Credit [3 points]** (this is an optional exercise) Implement seam carving:
 - (a) Compute magnitude of gradients of an image
 - (b) Find the connected path of pixels that has the smallest sum of gradients. A path is **valid** if it is connected (the neighboring points in the path are also neighboring pixels in the image), it starts in the first row of the image and in each step continues one row down. It finishes in the last row of the image.
 - (c) Remove the pixels in the path from the image. This gives you a new image with one column less.
 - (d) Remove a few paths with the lowest sum of gradients. Create (fun!) examples with a few images.
 - (e) Could you use this algorithm (with minor modifications) to find a skyline in `mountains.jpg`? A skyline is a connected path starting in the first column and finishing in the last column, separating the sky and the mountains.

You can find more details about seam carving in this paper:

S. Avidan and A. Shamir, *Seam Carving for Content-Aware Image Resizing*, SIGGRAPH 2007, <http://www.win.tue.nl/~wstahw/edu/2IV05/seamcarving.pdf>