

Intro to Image Understanding (CSC420)

Assignment 1

Posted: Sept 23, 2015 Submission Deadline : Oct 3, 11.59pm, 2015

Instructions for submission: Please write a document (either pdf, doc, etc) with your solutions (include pictures where needed). Include your code inside the document.

Max points: 12, max extra credit points: 3

- [1.5 points]** Write pseudo-code for computing convolution 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 convolution.
 - [1.5 points]** Is it possible to write a convolution in one pixel as a dot product between two vectors? Is it possible to write the full convolution between the image and the filter via matrix multiplication?
- [1 point]** Given a $n \times n$ image, I , and $m \times m$ filter, h , what is the computational cost of computing $h * I$ (the convolution)? What is the computational cost if h is a separable filter?
 - [1 point]** If I first convolve an image with a Gaussian filter with $\sigma = 1$, and then convolve the output with a Gaussian with $\sigma = 2$, this gives an equivalent result as if I just convolve the image with a Gaussian with what σ ?
 - [1 point]** Write a function that creates a Gaussian filter with σ as an input parameter.
 - [1 point]** Convolve the attached `image.png` with a (2D) Gaussian filter with $\sigma = 10$ and visualize the result (display the result of the convolution). You can use built-in functions for convolution.
 - [1 point]** Gaussian filter is separable. How can you use this fact to speed up convolution? What are the vertical and horizontal filters? No need to write code.
- [1 point]** Compute magnitude of gradients for the attached images `waldo.png` and `template.png`.
 - [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).
- [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`.
 - [1 point]** Any idea how you could find the bounds of the court in the image? No need to write code, just brainstorm!
- Extra Credit [3 points]** (this is an optional exercise) Implement seam carving:
 - 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 `toronto.jpg`? A skyline is a connected path starting in the first column and finishing in the last column, separating the sky and the buildings.