6.819/6.869: Advances in Computer Vision

Problem Set 1
(15 points)

Due: Thursday, October 1, 2015

Print a copy to give in class on Oct 1 and send an email to Aude (oliva@mit.edu) with the write-up and code files

What to submit: A short write-up and for some questions, MATLAB code. Please comment your code so we know which parts correspond to each question. Your write-up should include a very brief description of what your code is doing and a sample result on a test image, along with any other requested figures and the answers to any questions posed in the problem. PSet 1 covers lectures 2, 3 and 4 (Basics of Image Processing)

Q1(0.25 pt) Is there a kernel filter you can convolve with the image *g (left)* so as to obtain image f (right)? Justify why.

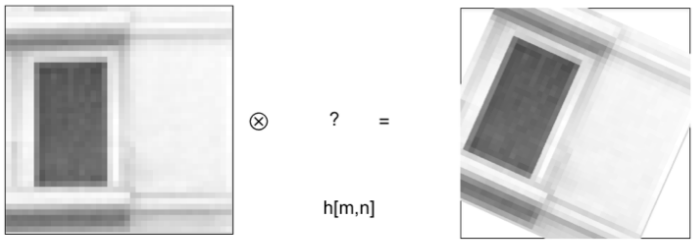

Q2- (0.25 pt) which kind of kernel h was used to create image f ? Draw it

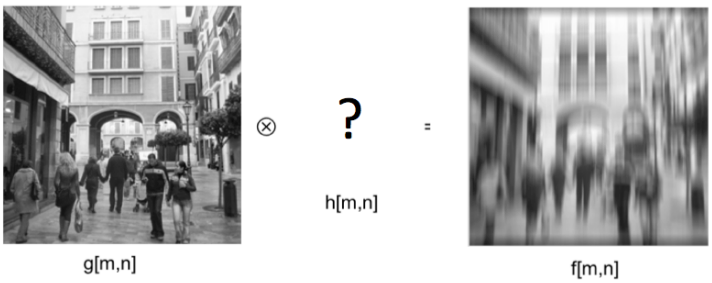

Q3 (1.5 pt). One application of a Laplacian operator is to restore fine details. Using a Laplacien kernel, perform image sharpening on the image "einsteinandwho.jpg" (provided). As answers, print the matlab code (including the kernel used) and the resulting sharpened image (note: the result should show the original image with finer details, not the filtered Laplacian image).

Q4 (1.5 pt): Sobel is a standard filter for edge detection. Using the image "mit.jpg" (provided).

(1) filter the image with an horizontal sobel
(2) filter the image with a vertical sobel
(3) combine a Gaussian (of your choice) and vertical sobel filter to get low spatial frequency edges

Print the matlab code and the three resulting images

Q5 (3.5 point): Use the hybrid image "einsteinandwho.jpg".

1) (*0.5 pt)* Who is the individual represented in the very low spatial frequency band of this Einstein hybrid?
2) *(3 pts)* Using the method or pyramid decomposition of your choice, remove the percept of the individual represented in the very low spatial frequency, so as to create an image of Einstein without the other hidden face.

Q6 (4 pts): The matlab code hybrid provided allows creating a hybrid image combining two different photos. Feel free to use another code

1) *(2 pts)* Create a hybrid image with a photo of yourself (your face) and the face of someone else. Or choose two photos of two individuals if you want to try to make a famous hybrid. It is advised to align first the features of the two faces for a better result. Print: the two original images; the hybrid; the low spatial frequency image, the high spatial frequency image; the two filters used and write a short the rational for the choices you made. Print a pyramid version (laplacian or steerable) of your hybrid.

2) *(2 pts)* Create a hybrid with three images of your choice: the hybrid should represent one image in the low spatial frequency band; another image in the high spatial frequency band; and a third one in the medium spatial frequency band. Print the 3 original images, the tri-hybrid, a version of the three filters used; and a pyramid version (laplacian or steerable) of your 3-hybrid.

Q7 (2 pts) - Image blending using a Laplacian Pyramid

Code for the Laplacian pyramid is available in the pyramid image processing toolbox (http://www.cns.nyu.edu/~eero/software.php)

Implement the function PyrBlend(im1,im2,mask) that takes as input two images and a binary mask (determining which part to use from each image) and produces the Laplacian pyramid blend of the two images. Use your function to blend two images of your favorite pets, friends or objects. Include in your report the original images, their Laplacian pyramids, the blending mask, and the resulting blended image.

Q8 (0.5 pt). **Fractals** are complex patterns that are self-similar across different scales. What do you expect the slope of the power spectrum along the 4 orientations (0, 45, 90 and 135 degrees) to be? You can answer the questions with a description, or collect ~ 20 fractals images and compute the results.

Q9 (1.5 pts). The power spectrum average of a set of images allows summarizing the statistics of orientations and spatial frequency ranges of the group. Select two groups of images of your choice from a category of places or objects or images (e.g. aerial images, paintings from particular epoch; etc). For the report, show a sample of 5 images per group; their image average; the average power spectrum, the slopes of 0,45,90 and 135 deg of the average power spectrum, and a short description of the results.

Extra points question:

Q10 (0.5 pt): What is special about the size *32 x 32* pixels (chosen by Torralba to create the 80 millions image dataset)?

Q11 (0.5 pt): What is bizarre about this image below?