Question 1.

Suppose you have an image I that contains an image of a left eye (the image is detailed enough that it makes a difference that it's the left eye).

- Write pseudocode to find other left eyes in the image using CC, NCC, SSD, etc.
- Using convolution (rather than cross-correlation, write pseudocode to find other left eyes in the image
- If we simply filter the image using an image of a left eye using cross-correlation and look for maxima in the results, what kind of images of left eye won't we find? How to address this issue?

Question 2.

Suppose now you have an image of the UofT crest **lying on a white field**. What issue arises with trying to use SSD etc. to find where the UofT crest is in the image? How to resolve this issue?

Question 3.

Give an example of a filter for which filtering using cross-correlation and filtering using convolution produces the same results. Give an example of a filter for which filtering using cross-correlation and filtering using convolution produces different results.

Question 4.

Suppose we are trying to find eyes in an image by filtering the image with a filter that looks like an eye, without subtracting the mean of the window in the image. We then look for maxima in the results. Why would this generally produce false detections?

Question 5.

Suppose we are trying to find eyes in an images by filtering the image with a filter that looks like an eye, without subtracting the mean of the window of the image. We then look for maxima in the results. Why would this generally produce false detections?

Question 6.

- Write pseudo-code to half-size an image. Do not use imresize() or similar functions
- Same as above, but make sure to avoid aliasing
- Give an example of an image (for example, represent it as a small 2D array) where subsampling will produce aliasing

Question 7.

- Write pseudo-code to half-size an image. Do not use imresize() or similar functions
- Same as above, but make sure to avoid aliasing

- Given an example of an image (for example, represent it as a small 2D array) where subsampling will produce aliasing

Question 8.

Write pseudocode which transforms a given grayscale image so that the shapes in the image are preserved, but the minimum intensity of the image is 0.4, and the maximum intensity of the image is 0.6

Question 9.

Write pseudocode that smoothes an input image

Question 10.

- Write pseudocode to smooth an image using Guassian smoothing and mean-filter smoothing
- Why is Gaussian smoothing generally preferable to mean-filter smoothing?

Question 11.

- Give an example of a filter such that if we filter with it, the image does not change
- Give another example of such a filter

Question 12.

- Give an example of a filter such that if we filter with it, the image does not change
- Give another example of such a filter

Question 13.

What is a filter that, if you convolve the image on the left with with it, will produce output whose absolute value is shown on the right?

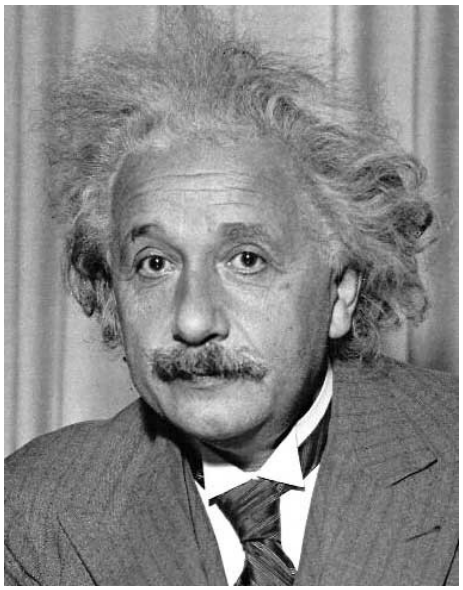
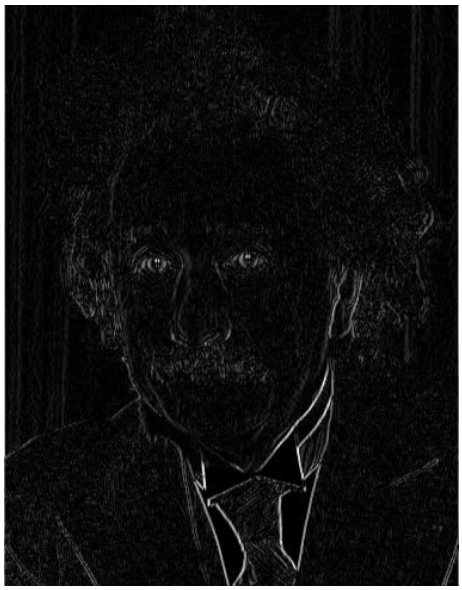

Question 14.

Suppose you compute the Gaussian pyramid of a 256×256 image by halfsizing the image repeatedly. Each pixel takes one byte to store. What is the amount of memory needed to store the entire Gaussian pyramid?

Question 15.

What is the result of applying median filtering to the 1D image `[10 15 20 23 27 30 31 33 90]`? The size of the median filter is 3.

Question 16.

Write pseudocode to add salt-and-pepper noise to an image.

Why is median filtering more effective than Gaussian filtering for eliminating salt-and-pepper noise?

Question 17.

Why is median filtering more effective than Gaussian filtering for eliminating salt-and-pepper noise?

Question 18

What is the gradient of a function? Explain how to approximate the gradient of an image, and explain the relationship between the gradient of a function in calculus and the gradient of an image

Question 18.

Name 4 factors that cause edges in images.

Question 19.

Explain the relationship between the image gradient and edges in the image, and explain why this relationship holds.

Question 20.

Explain why Gaussian noise makes the image Gradient very noisy if it is computed without smoothing the image first. Provide (small!) examples of a 1-dimensional image and a 1-dimensional image corrupted with Gaussian noise (that you made up) to illustrate your explanation.

Question 21.

Explain why smoothing solves the problem in the question above.

Question 22.

If the edge normal is $(2, 3)$, what is the direction of the edge? What is the local image contrast (i.e., the difference between pixel intensities for pixels that lie on different sides of the edge)

Question 23.

If the image gradient is $(-1, 1)$, what is the magnitude of the gradient? What is the orientation of the gradient (in degrees)

Question 24.

Write pseudocode to perform non-maximum suppression. What are the inputs that you need?

Question 25.

Explain why using a single threshold on the magnitude of the gradient to detect edges is sometimes insufficient.

Question 26.

The edge images below were produced using different σ when smoothing before detecting Canny edges. For which image is the σ larger? Why?



Question 27.

What kind of boundaries inside the image does Canny edge detection in images which were converted to grayscale not detect?

Question 28.

Explain bilinear interpolation, provide a numerical example of performing bilinear interpolation.

Question 29.

What kind of artifacts does upsampling using nearest-neighbours produce?

Question 30.

Suppose you are upsampling a 1D image by a factor of 2. What filter (explicitly give the mask, do not just draw a picture) would produce a linear interpolation when applied to the upsampled image where the "new" pixels are initially set to 0?

Question 31.

What factors make face detection difficult?

Question 32.

Explain how one can denoise images of particular objects using PCA. Provide pseudocode.

Question 33.

If `V` is a matrix whose columns are the principal components of centered faces, `x` is a grayscale image (represented as a 2D array), and `mean_face` is an image of the average face (a 2D array). Write Python

code that computes the reconstruction of `x` in the vector space spanned by the first `20` principal components stored in `V`

Question 34.

- Explain how one might compress images of particular objects using PCA.
- Bob claims that compression using PCA makes no sense, since if you compress e.g. a face using PCA, you need to transmit all of the average face, the principal components of the face space, and the corresponding coefficients for the principal components to the recipient so that the recipient can reconstruct the face, and that would involve transmitting more data than simply transmitting the intensity of each pixel. Why is this argument incorrect?

Question 35.

Write Python code to generate the images of the gratings in slide 27 of `FreqDomain.pdf` (assume that you have to work out what `k` and `l` might be.)

Question 36.

Write Python code to draw a line that's parallel to the unit vector (u, v) and that passes through the centre of an image.

Question 37.

For the gratings given on slide 27 of `FreqDomain.pdf`, work out the `k` and `l` values. Assume you know the size of the images (e.g., 100 by 200).

Question 38.

For the Fourier Images on slide 29 of `FreqDomain.pdf`, work out the intensity images that correspond to them. Assume the size of the images is 100×100

Question 39.

For the Fourier Image on slide 32 of `FreqDomain.pdf`, identify the things in the original image that correspond to each of the lines in the Fourier Image. (Which edge corresponds to which line?)

Question 40.

For slides 37 and 38 of `FreqDomain.pdf`, given the two intensity images, identify what was done to the Fourier Image.

Question 41.

Write code to perform Unsharp Masking of an image.