

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 Gaussian 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 \mathbf{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.

Question 42.

Two images, I_1 and I_2 , are of the same size. Write code to combine those images into one images of the same size as I_1 , with I_1 transitioning into I_2 midway through the width of the image. (See slide 5 of [blending.pdf](#)). Use feathering.

Question 43.

What are the issues with using feathering with a transition window of size w for blending? How to select the size of the window used in feathering to avoid those issues? Explain. When is it not possible to select a good window size?

Question 44.

Write code to produce the k -th level of the Laplacian pyramid of image I .

Question 45.

Suppose $L[i]$ is the i -th level of the Laplacian pyramid. There are 5 levels in total, with $L[4]$ being the last level, which is just the 4-th level of the Gaussian pyramid of the image. Write Python code to obtain the original image.

Question 46.

Suppose you are given a database with faces, and a name is given for each face. Suppose you are given a new input face. Outline the algorithm to recognize that face using the Eigenface decomposition of the database of faces.

Question 47.

What are the limitations of doing face recognition with Eigenfaces? Suppose you're given a database of people with bounding boxes around the faces. What must be true about the bounding boxes for face recognition with Eigenfaces to work (be more specific than "the bounding boxes must be accurate")?

Question 48.

How would you determine how many eigenfaces to use for face recognition?

Question 49.

Suppose you have a composite image C , and an image B . Write code to change the background in C to be B using the Known Background Matting technique. Make any assumptions and introduce any variables that you need.

Question 50.

Suppose you have a composite image C , and an image B . Write code to change the background in C to be B using the Blue Screen Matting technique. Make any assumptions and introduce any variables that you need.

Question 51.

What are the downsides of Known Background Matting?

Question 52.

What are the downsides of Blue Screen Matting?

Question 53.

Derive the equations on slide 18 of `matting.pdf` (the system of equations for grey objects). Derive a similar set of equations for skin-coloured foreground objects.

Question 54.

Explain how triangulation matting works. Derive the equations on slide 20.

Question 55.

Write code to produce a mirror reflection of image I , followed by a rotation by 42° around $(1, 1)$

Question 56.

Given the code in `homog.py` on the course webpage, modify it to rotate an image using bilinear interpolation

Question 57.

What is the transformation matrix that corresponds to a rotation by 30° around the origin followed by a translation, followed by a magnification by a factor of 2?

Question 58.

Write code to rotate image I around its centre by 30° clockwise. Use nearest-neighbour interpolation.

Question 59.

How many correspondences are required to recover a Euclidian (translation+rotation) transformation. Why? (If you answer in terms of DOF, explain how DOF related to the number of equations required)

Question 60.

Outline how image morphing works

Question 61.

Suppose we know that the pixel with the colour (200, 220, 255) in the image is actually a part of a white object. How could that have occurred? (Answer: e.g., the illumination wasn't white). Propose a method to white-balance the image.

Question 62.

Suppose an intensity in the range of 0..10000 is given for an image, and most of the intensity are close to 0. How would one transform the image so that it could be displayed well on a display that supports the 0..255 range?

Question 63.

Draw typical intensity histograms for low-contrast images and high-contrast images. Draw typical histograms for dark images and bright images.

Question 64.

Draw a typical intensity histogram for a histogram-equalized image. What is histogram equalization useful for?

Question 65.

Give an example of an image (e.g., a 2×2 image) which would not be affected by histogram equalization.

Question 66.

Write Python code to filter an image using a bilateral filter (the details are up to you). When removing noise, what is the advantage of filtering using a bilateral filter over filtering using a gaussian filter?

Question 67.

In flash/no-flash photography, we use a *joint bilateral filter* that uses the information from both photos. Explain how it works.

Question 68.

Given the formula on the last slide of `flashnoflash.pdf`, explain what each term does. What does the `argmin` part there mean?

Question 69.

Given two triangles, and assuming you know which vertices correspond to which, explain precisely (i.e, you should be able to code this up) how to warp one triangle onto another triangle.

Question 70.

Explain how Hybrid Images work

Question 71.

Explain how data-driven beautification works (slides 3-4 are enough)