

Duration: **105 minutes**
 Aids Allowed: **Non-programmable calculator**

Student Number: _____

Family Name(s): _____

Given Name(s): _____

Lecture Section: Afternoon Section

Evening Section

*Do **not** turn this page until you have received the signal to start.
 In the meantime, please read the instructions below carefully.*

This term test consists of 6 questions on 20 pages (including this one), printed on both sides of the paper. *When you receive the signal to start, please make sure that your copy of the test is complete, fill in the identification section above, and write your name on the back of the last page.*

Answer each question directly on the test paper, in the space provided, and use the reverse side of the pages for rough work. If you need more space for one of your solutions, use the reverse side of a page and *indicate clearly the part of your work that should be marked.*

You will not lose marks for minor syntax errors. If you cannot complete a helper function that is required to solve a problem, you should clearly indicate what the function that you cannot complete is supposed to do. You will receive part marks as appropriate for knowing what kind of helper function is needed and for using it correctly.

Write up your solutions carefully! Comments and docstrings are *not* required to receive full marks. However, they may help us mark your answers, and part marks *will* be given for showing that you know the general structure of an answer, even if your solution is incomplete.

MARKING GUIDE

1: _____/ 6

2: _____/ 5

3: _____/ 4

4: _____/ 6

5: _____/ 6

6: _____/ 1

TOTAL: _____/28

Use this page for rough work—clearly indicate any section(s) to be marked.

Question 1. [6 MARKS]

Suppose you have 3 grayscale images I1, I2, and I3, all of the same size. You would like to determine whether I1 is more similar to I2 or I3.

Part (a) [2 MARKS]

Describe a situation where comparing the images using Zero-Mean Normalized Cross-Correlation is preferable to comparing them using the Sum of Squared Differences between them.

Part (b) [2 MARKS]

Write Python code that outputs “1~2” if I1 is more similar to I2 than to I3, “1~3” if I1 is more similar to I3 than to I2, and “tie” otherwise. **Compare the images using SSD.** You may use standard NumPy functions.

Use this page for rough work—clearly indicate any section(s) to be marked.

Part (c) [2 MARKS]

Write Python code that outputs “1~2” if I1 is more similar to I2 than to I3, “1~3” if I1 is more similar to I3 than to I2, and “tie” otherwise. **Compare the images using Zero-Mean NCC.** You may use standard NumPy functions.

Use this page for rough work—clearly indicate any section(s) to be marked.

Question 2. [5 MARKS]

Consider the following image. The image intensity at each pixel is indicated inside the pixel, and the point with coordinates $(1.8, 1.8)$ is marked. Suppose you need to resize the image to make it larger, and so you need an estimate of the intensity in $(1.8, 1.8)$, which you want to obtain using interpolation.

10	10	10	10	10
10	20 <small>$(1.8, 1.8)$•</small>	10	10	10
10	30	40	10	10
10	10	10	10	10
10	10	10	10	10

Part (a) [2 MARKS]

Estimate the intensity at point $(1.8, 1.8)$ using **nearest-neighbour interpolation**. Briefly explain how you got your answer.

Use this page for rough work—clearly indicate any section(s) to be marked.

Part (b) [3 MARKS]

Estimate the intensity at coordinates $(1.8, 1.8)$ using **bilinear interpolation**. Show your work.

Use this page for rough work—clearly indicate any section(s) to be marked.

Question 3. [4 MARKS]

Consider the image at the bottom of the page again.

Part (a) [2 MARKS]

Compute (i.e., your answer should be numerical) the gradient vector at the centre of pixel (2,1) (Note: the intensity of that pixel is 30). Any reasonable method is acceptable. Show your work. You do not need to perform Gaussian smoothing. Keep in mind that in the image, pixel (0,0) is in the top-left corner.

Part (b) [2 MARKS]

Draw a vector on the image, starting in pixel (2,1), which points in approximately the same direction as the gradient at (2,1).

10	10	10	10	10
10	20 <small>(1.8, 1.8)•</small>	10	10	10
10	30	40	10	10
10	10	10	10	10
10	10	10	10	10

Use this page for rough work—clearly indicate any section(s) to be marked.

Question 4. [6 MARKS]**Part (a)** [3 MARKS]

Suppose V is a matrix whose rows are the Principal Components of centered faces (sorted by their corresponding eigenvalue of the covariance matrix, with the 0 – th row corresponding to the largest eigenvalue), 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 squared reconstruction error of x in the vector space spanned by the first 20 principal components stored in V .

Part (b) [3 MARKS]

What can you say about x if the squared reconstruction error computed above is close to 0?

Use this page for rough work—clearly indicate any section(s) to be marked.

Question 5. [6 MARKS]**Part (a)** [2 MARKS]

Suppose h is an image of size $N \times M$ and suppose that \hat{h} is the FT of image h . If $\hat{h}(0,0) = 1$ and $\hat{h}(k,l) = 0$ for $(k,l) \neq (0,0)$. Give the most precise possible description of image h . Justify your answer.

Part (b) [4 MARKS]

Suppose image I is filtered using an unsharp mask to produce a sharpened image S . State precisely what the relationship is between the Fourier Spectrum image of I and the Fourier Spectrum image of S . Justify your answer.

Use this page for rough work—clearly indicate any section(s) to be marked.


Question 6. [1 MARK]

Consider the photo of Robarts Library below (scale is shown). The roof of the building is pictured at an angle. In Fourier Domain, the edge between the roof and the sky (and surrounding image elements) is represented by a line of high coefficients in the power spectrum. Draw this line in the coordinate system below. Be as precise as possible. Explain how you got your answer.


Reminder: the 2D Discrete Fourier Transform of an image h of size $N \times M$ is:

$$\hat{h}(k, l) = \sum_{n,m} h(n, m) \exp(-i(\omega_k n + \omega_l m)).$$

Here, $\omega_k = (2\pi k)/N, \omega_l = (2\pi l)/M$.



Image



Fourier Spectrum (with coordinate grid)

Use this page for rough work—clearly indicate any section(s) to be marked.

Additional page for answers

On this page, please write nothing except your name.

Family Name(s): _____

Given Name(s): _____