ENGINEERING TRIPOS PART IIB

Thursday 12 May 2011 2.30 to 4

Module 4F8

IMAGE PROCESSING AND IMAGE CODING

Answer not more than ***three*** *questions.*

All questions carry the same number of marks.

The ***approximate*** *percentage of marks allocated to each part of a question is indicated in the right margin.*

There are no attachments.

STATIONERY REQUIREMENTS
Single-sided script paper

SPECIAL REQUIREMENTS
Engineering Data Book
CUED approved calculator allowed

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

JL03

1 (a) The inverse Fourier transform of some desired zero-phase 2-D frequency response will not normally give an impulse response with finite support.

(i) Explain how the *windowing method* is used to create finite support filters and describe the effect of windowing on the ideal frequency response. Discuss two methods of forming 2-D window functions from 1-D window functions. [15%]

(ii) Consider the following 1-D triangular window functions:

$$w_k(u_k) = \begin{cases} 1 - \frac{|u_k|}{U_k} & \text{if } |u_k| < U_k \\ 0 & \text{otherwise} \end{cases}$$

for $k = 1, 2$. Find the form of the spectrum of the 2-D window function formed via the product of the above window functions, $w_1(u_1)\,w_2(u_2)$. [30%]

(iii) Comment on the properties of this spectrum in terms of its effects on a given ideal frequency response. [15%]

(b) Consider the zero-phase ideal frequency response $H(\omega_1, \omega_2)$ shown in Fig. 1, where H takes the value 1 in the shaded region and zero outside this region.

(i) If an image is sampled with spacings Δ_1 and Δ_2 in the u_1 and u_2 directions respectively, find the ideal impulse response, $h(u_1, u_2)$, corresponding to $H(\omega_1, \omega_2)$ (you may use standard results). Verify that if $\Omega_U = \Omega_L$, your solution reduces to the response of a simple low pass filter. [30%]

(ii) By looking at the behaviour of this ideal impulse response along the axes and along the diagonals, describe and give a rough sketch of its form. [10%]

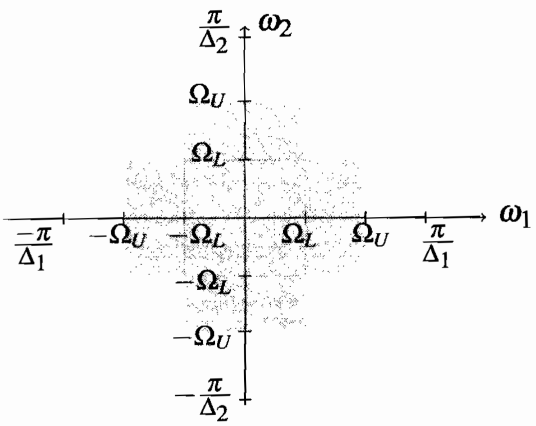

Fig. 1

2 (a) Assume that an observed image, $y(u_1,u_2)$ can be modelled as a linear distortion of the true image, $x(u_1,u_2)$, plus additive noise, $d(u_1,u_2)$.

(i) Explain how we might find an empirical estimate of the distorting function, L, for a given imaging system. [10%]

(ii) If the distorting function is known and the noise is neglected, explain what is meant by *inverse filtering* and describe how this is used to estimate the true image. Explain why such estimates are often poor for real images and describe how the *generalised inverse filter* can improve performance. [20%]

(iii) The Wiener filter is a linear filter which exhibits improved performance over inverse filtering. The Wiener filter can be derived via a Bayesian approach in which we estimate the posterior, $P(\mathbf{x}|\mathbf{y})$, using Bayes theorem. In the following expression for $P(\mathbf{x}|\mathbf{y})$ for the Wiener filter,

$$P(\mathbf{x}|\mathbf{y}) \propto \mathrm{e}^{-\frac{1}{2}[(\mathbf{y}-L\mathbf{x})^T N^{-1}(\mathbf{y}-L\mathbf{x})+\mathbf{x}^T C^{-1}\mathbf{x}]}$$

describe the nature of the matrices N and C, noting any assumptions that are made and give the form of the prior probability used. Outline how, by taking alternative priors for $\mathbf{x}$, we are able to improve on the Wiener filter and give an example of one such prior. [30%]

(b) Suppose we are given an $m \times n$ image $g(n_1,n_2)$, which is a sampled version of the continuous image $g(u_1,u_2)$. The 2-D Fourier transform of g is $G(\omega_1,\omega_2)$, which is zero for the ranges $|\omega_1| > \Omega_1$ and $|\omega_2| > \Omega_2$.

(i) You may assume that the sampling at the original resolution ($m \times n$) displays no signs of aliasing. In order to minimise storage space, we would like to store downsampled versions of images. Find the minimum size, (m_{min}, n_{min}), to which the image g can be reduced in order to ensure no aliasing effects will be present. [30%]

(ii) Explain how the distortion effects, which are characteristic of aliasing in images, occur. [10%]

(TURN OVER

3 The elements of a one-dimensional (1-D) discrete cosine transform (DCT) matrix T for n even are given by:

$$t_{1i} = \sqrt{\frac{1}{n}} \qquad \text{for } i = 1,..,n,$$

$$t_{ki} = \sqrt{\frac{2}{n}} \cos\left(\frac{\pi(2i-1)(k-1)}{2n}\right) \qquad \text{for } i = 1,..,n, \; k = 2,..,n.$$

(a) Explain how a 1-D transform matrix T (of size $n \times n$) may be used to perform a 2-D transformation of an $n \times n$ block of image pixels X into an $n \times n$ block of transform coefficients Y. [15%]

(b) Considering the case of $n = 8$, show how the symmetries present in the rows of the transform matrix T allow an efficient strategy to be developed for computing the product $T\mathbf{x}$ where $\mathbf{x}$ is an n-element column vector. Calculate the amount by which the computation may be reduced, compared with direct multiplication of $\mathbf{x}$ by T, in terms of numbers of multiplies and numbers of add/subtract operations. How does this result scale up for the 2-D transform of an 8×8 block of pixels? [40%]

(c) The subimages, formed from the rearranged transform coefficients obtained by applying an 8×8 2-D DCT to an image of size 1024×768 pixels, are each of size 128×96 and are labelled $U_{i,j}$ for $i = 1,..,8$ and $j = 1,..,8$. The mean entropy of the coefficients in each subimage is given approximately by:

$$H_{i,j} = \frac{6}{i+j-1} \quad \text{bits / coefficient.}$$

Estimate the number of bits required to encode this image. [25%]

(d) Now considering the cases $n = 2, 4, 8$ and 16, briefly discuss the relative merits of using these different sizes of transform for image compression. [20%]

4 (a) Draw the basic 2-band analysis filter bank with down-samplers, upon which wavelet transforms are based. Also draw the corresponding 2-band reconstruction filter bank with up-samplers. What is the perfect reconstruction requirement and why is it important? [25%]

(b) Draw a block diagram to show how a number of such analysis filter banks would normally be employed to create a 3-level two-dimensional wavelet transform which produces 10 output subbands. [25%]

(c) Using a test image of a light-coloured square on a dark background, explain what features of the square are picked out by each subband. Hence discuss how the wavelet transform decomposes a typical real-world image and why it results in energy compression of the image data. (Hint: it may be helpful to sketch an image of the wavelet coefficients in the various subbands and assume the simple Haar wavelet is used). [25%]

(d) Explain how reconstruction filter banks would be used to invert the wavelet transform of part (b) and reconstruct an approximation to the original image. Discuss what features the reconstruction filters should possess in order that the artifacts in the output image, caused by data compression, should have low visibility to the human observer. [25%]

END OF PAPER