Topics: Continuous 1 and 2D Fourier Transform

Spring 2009 Final: Problem 1 (CSFT and DTFT properties)

Derive each of the following properties.

a) Show that if $g(t)$ has a CTFT of $G(f)$, then $g(t-a)$ has a CTFT of $e^{-2\pi jaf}G(f)$.

b) Show that if $g(t)$ has a CTFT of $G(f)$, then $g(t/a)$ has a CTFT of $|a|G(af)$.

c) Show that if x_n has a DTFT of $X(e^{j\omega})$, then $(-1)^n x_n$ has a DTFT of $X(e^{j(\omega-\pi)})$.

d) Show that if $g\left(\begin{bmatrix} x \\ y \end{bmatrix}\right)$ has a CSFT of $G\left(\begin{bmatrix} u \\ v \end{bmatrix}\right)$, then $g\left(A\begin{bmatrix} x \\ y \end{bmatrix}\right)$ has a CSFT of $|A|^{-1}G\left(\left(A^{-1}\right)^t\begin{bmatrix} u \\ v \end{bmatrix}\right)$.

(Hint: Use the notation $r = \begin{bmatrix} x \\ y \end{bmatrix}$ and $f = \begin{bmatrix} u \\ v \end{bmatrix}$, so that $G(f) = \int_{\Re^2} g(r)e^{-jr^t f}dr$.)

Spring 2008 Exam 1: Problem 1 (CSFT)

Consider the CSFT given by

$$F(u,v) = \frac{1}{2}\left[\delta(u-u_o, v-v_o) + \delta(u+u_o, v+v_o)\right]$$

a) Calculate, $f(x,y)$, the inverse CSFT of $F(u,v)$.

b) Calculate the minimum distance between nearest peaks in the function $f(x,y)$.

c) Sketch $F(u,v)$ and $f(x,y)$ when $u_o = 4$ and $v_o = 3$. Label the axis on your sketch, and also make sure to label important dimensions of the signal and its transform.

Spring 2006 Exam 1: Problem 2 (CSFT)

a) Calculate the CSFT of

$$f(x,y) = rect\left(x/A, x/B\right)$$

b) Calculate the CSFT of

$$g(x,y) = \sum_{k=-\infty}^{\infty}\sum_{l=-\infty}^{\infty} f\left(x - 5kA, y - 5lB\right)$$

c) Calculate the CSFT of

$$h(x,y) = \begin{cases} g(x,y) & \text{for } |x| < T/2 \text{ and } |y| < T/2 \\ 0 & \text{otherwise} \end{cases}$$

d) Sketch $h(x,y)$ for $A = B = 1$ and $T = 50$.

Topics: Tomography and MRI

Spring 2008 Exam 1: Problem 4 (tomography)

Consider an X-ray imaging system shown in the figure below.

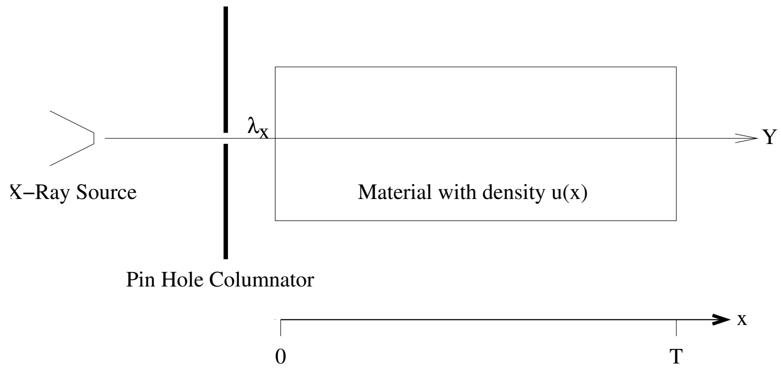

Photons are emitted from an X-ray source and columnated by a pin hole in a lead shield. The columated X-rays then pass in a straight line through an object of length T with density $u(x)$ where x is the depth into the object. The number of photons in the beam at depth x is denoted by the Poisson random variable Y_x with $E[Y_x] = \lambda_x$ where all distances are measured in units of cm and all absorption constants are measured in units of cm^{-1}.

a) Write a differential equation which describes the behavior of λ_x as a function of x.

b) Solve the differential equation of part a)

c) Specify how the photon counts Y_x can be used to compute the path integral of $u(x)$.

Spring 2008 Final: Problem 5 (MRI)

Consider an MRI that only images in one dimension, x. So for example, the object being imaged might be a thin rod oriented along the x-dimension.

In this example, assume that the magnetic field strength at each location is given by

$$M_o + G(t)x$$

where M_o is the static magnetic field strength and $G(t)x$ is the linear gradient field in the x dimension. Then the frequency of precession for a hydrogen atom (in rad/sec) is given by the product of γ, the gyromagnetic constant, and the magnetic field strength.

a) Calculate $\omega(x,t)$, the frequency of precession of a hydrogen atom at location x and time t.

b) Calculate $\phi(x,t)$, the phase of precession of a hydrogen atom at location x and time t assuming that $\phi(x,0) = 0$.

c) Calculate $r(x,t)$, the signal radiated from hydrogen atoms in the interval $[x, x+dx]$ at time t.

d) Calculate $r(t)$, the signal radiated from hydrogen atoms along the entire object.

e) Calculate an expression for $a(x)$, the quantity of processing hydrogen atoms along the thin rod, from the function $r(t)$.

Spring 2008 Exam 1: Problem 2 (Tomography)

Assume that we know (or can measure) the function

$$p(x) = \int_{-\infty}^{\infty} f(x,y)dy \ .$$

Using the definitions of the Fourier transform, derive an expression for $F(u,0)$ in terms of the function $p(x)$.

Topics: Discrete transforms; 1 and 2D Filters, sampling, and scanning

Spring 2010 Exam 2: Problem 3 (sampling)

Consider a sampling system were the input, $s(t) = \text{sinc}(2t)$, is sampled with period $T = 1/2$ to form the sampled signal $x(n) = s(nT)$.

After sampling, you determine that you selected the wrong sampling rate, and really need to have sampled the signal at the period $T_2 = 1/4$; so you interpolate by a factor of $L = 2$ to form the signal $y(n)$.

a) Sketch the signal $s(t)$ and its CTFT $S(f)$. What is the Nyquist sampling rate for this signal?

b) Sketch the signal $x(n)$ and also sketch its DTFT $X(e^{j\omega})$.

c) Sketch $x(n)$ after it is up-sampled by $L = 2$.

d) Sketch the interpolation filter's impulse response.

e) Sketch the signal $y(n)$.

f) What is the relationship between $y(n)$ and $s(t)$?

Spring 2009 Exam 1: Problem 1 (FIR filters and frequency response)

Consider the linear-space invariant FIR filter given by

$$y(m, n) = x(m, n) * h(m, n)$$

where

$$h(m, n) = \begin{cases} \pi & \text{if } |m| \leq 5 \text{ and } |n| \leq 5 \\ 0 & \text{otherwise} \end{cases}$$

a) Sketch the function $h(m, n)$. You may use any method you prefer to sketch it (i.e. 2D or 3D sketch), but make sure to clearly show the zero and nonzero values and their locations in the plane.

b) Calculate $H(e^0, e^0)$, the DC gain of the FIR filter.

c) Is this function separable? If so, then give its separable decomposition $h(m, n) = g(m)f(n)$.

d) How many multiplies per output point are required for direct implementation of the FIR filter output?

e) Specify an alternative implementation which uses the separable nature of the FIR filter.

f) How many multiplies per output point are required for separable implementation of the FIR filter output?

Spring 2007 Exam 1: Problem 1 (DSFT and 2D Z-transform)

Consider the following 2D system with input $x(m,n)$ and output $y(m,n)$.

$$y(m,n) = x(m,n) + \lambda \left(x(m,n) - \frac{1}{9} \sum_{k=-1}^{1} \sum_{l=-1}^{1} x(m-k, n-l) \right) .$$

a) Is this a linear system? Is this a space invariant system?

b) What is the 2D impulse response of this system, $h(m,n)$?

c) Calculate the frequency response, $H(e^{j\mu}, e^{j\nu})$?

d) Describe how the filter behaves when λ is positive and large.

e) Describe how the filter behaves when λ is negative and > -1.

Spring 2007 Exam 1: Problem 2 (DSFT and 2D Z-transform)

Consider the causal linear space invariant system with input $x(m,n)$ and output $y(m,n)$ that is specified by

$$y(m,n) = x(m,n) + ay(m-1,n) + by(m,n-1)$$

a) Calculate the transfer function $H(z_1, z_2)$ for this system.

b) Calculate the value of

$$\sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} h(m,n)$$

where $h(m,n)$ is the 2D impulse response of the system.

c) Calculated the value of

$$\sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} h(m,n) \cos(\omega_o m) .$$

d) Is this system stable for all, none, or some values of (a,b)? Justify your answer.

Topics: Random processes, spectral estimation, and eigen-image analysis

Spring 2010 Final: Problem 5 (power spectrum and MMSE prediction)

Let $X = [x_1, \cdots, x_N]$ be a $P \times N$ matrix formed by P dimensional column vectors, $x_n \in \Re^P$ where $N < P$. We will assume that each column vector is an independent multivariate Gaussian random vector with distribution $N(0, R)$, where R is a positive definite and symmetric matrix.

Furthermore, let $X = U\Sigma V^t$ be the singular value decomposition of X, where $\Sigma_{i,i} \geq \Sigma_{j,j}$ when $i < j$.

a) Write a simple matrix expression for the sample covariance, $\hat{R}$.

b) Let $\hat{R} = E\Lambda E^t$ be the eigen decomposition of the sample covariance matrix, where E is the orthonormal transform with eigenvectors as columns and Λ is the diagonal matrix of eigen values. (Without loss of generality assume that the eigenvalues are ordered from largest to smallest so that $\Lambda_{i,i} \geq \Lambda_{j,j}$ when $i < j$.

How many non-zero eigenvalues does the matrix $\hat{R}$ contain?

c) Specify the eigenvectors and eigenvalues of $\hat{R}$ in terms of the SVD of X.

d) In some application, you can only use two numbers to specify each vector, x_n. So each vector must be approximated by

$$x_n \approx a_n e_1 + b_n e_2$$

where $e_1 \in \Re^P$ and $e_2 \in \Re^P$ are two orthonormal vectors, and a_n and b_n are two scalar values used to specify each vector, x_n.

What is the best choice of e_1 and e_2?

Spring 2010 Exam 1: Problem 3 (power spectrum and MMSE prediction)

Let y_n be a wide-sense stationary, jointly Gaussian, zero-mean, discrete-time random process. Then we know that the minimum mean squared error (MSEE) predictor has the form

$$\begin{aligned} \hat{y}_n &= E\left[y_n | \, y_k \text{ for } k < n\right] \\ &= \sum_{i=1}^{\infty} h_{n,i} y_{n-i} \end{aligned}$$

for some scalar constants $h_{n,i}$.

a) How are the functions $h_{n,i}$ and $h_{k,i}$ related for all n, k, and i? Provide a precise justification for your answer.

b) Consider the function $z_n = y_{-n}$. What is the MMSE predictor for z_n? Provide a precise justification for your answer.

c) What is the autocorrelation of $x_n = y_n - \hat{y}_n$. Provide a precise justification for your answer.

d) Derive an expression for the power spectrum of the random process, y_n.

Spring 2010 Exam 1: Problem 1 (power spectrum and IIR filters)

Consider the following 2-D discrete-time linear system.

$$y(m,n) = x(m,n) + ay(m-1,n) + by(m,n-1) - aby(m-1,n-1)$$

where a and b are scalar constants.

a) Compute the transfer function $H(z_1, z_2)$ for the system.

b) Compute the impulse response $h(m,n)$ for the system.

c) For what values of a and b is the system stable.

d) Compute $S_y(e^{j\mu}, e^{j\nu})$, the power spectrum of $y(m,n)$, when $x(m,n)$ is a set of i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables.

Topics: Neighborhoods, connected components, clustering, and edge detection

Spring 2010 Exam 2: Problem 1 (edge detection)

Your objective is to perform edge detection on the sampled image $g(m,n) = f(mT, nT)$, where $f(x,y)$ is the associated continuous space image and $T = 1$. You will do this using a combination of gradient and Laplacian based operators.

a) Specify the condition for the detection of edges on the continuous image $f(x,y)$ using derivatives over x and y, and a single threshold γ.

b) Specify an approximate discretized gradient operator for the image $g(m,n)$.

c) Specify an approximate discretized Laplacian operator for the image $g(m,n)$.

d) Specify the condition for the detection of edges on the discretized image $g(m,n)$ using approximate discretized gradient and Laplacian operators.

e) Describe how the threshold γ should be selected. What are the tradeoffs in its selection?

Spring 2007 Exam 2: Problem 3 (edge detection)

Consider the linear time-invariant discrete-time filter

$$y(n) = x(n) * h(n)$$

with input $x(n)$, output $y(n)$, and impulse response $h(n)$. Further, assume that $x(n)$ is created by sampling a continuous-time signal $s(t)$ as

$$x(n) = s(nT)$$

where $T = 1$.

a) Specify a simple FIR filter $h(n)$ so that $y(n)$ is approximately equal to $\left.\frac{ds(t)}{dt}\right|_{t=n-\frac{1}{2}}$.

b) Specify a simple FIR filter $h(n)$ so that $y(n)$ is approximately equal to $\left.\frac{ds(t)}{dt}\right|_{t=n+\frac{1}{2}}$.

c) Specify a simple FIR filter $h(n)$ so that $y(n)$ is approximately equal to $\left.\frac{d^2s(t)}{dt^2}\right|_{t=n}$.

d) Specify an operation on $y(n)$ which determines when $\frac{d^2s(t)}{dt^2} = 0$ for some value of $n \leq t \leq n+1$.

Spring 2004 Midterm Exam: Problem 2 (connected components)

Consider the following main program and subroutine.

Main Routine:
$ClassLabel = 1$
Initialize $Y_r = 0$ for $r \in S$
For each $s \in S$ in raster order {
  if($Y_s = 0$) {
    ConnectedSet($s, Y, ClassLabel$)
    $ClassLabel \leftarrow ClassLabel + 1$
  }
}

Subroutine:
ConnectedSet($s_0, Y, ClassLabel$) {
  $B \leftarrow \{s_0\}$
  While B is not empty {
    $s \leftarrow$ any element of B
    $B \leftarrow B - \{s\}$
    $Y_s \leftarrow ClassLabel$
    $B \leftarrow B \bigcup \{r : r \in c(s) \text{ and } Y_r = 0\}$
  }
  return(Y)
}

Also consider the following binary image

0	1	0	0	1
1	0	0	1	1
0	1	1	0	0
0	1	1	0	0
0	1	0	0	1

a) Calculate the output when the binary image is process by the main routine using a 4-pt neighborhood. Write your result in the table below.[1]

b) Calculate the output when the binary image is process by the main routine using an 8-pt neighborhood. Write your result in the table below.[2]

[1] Pixels on the image edge should be consider to have only 3 neighbors, and pixels in image corners should be considered to have only 2 neighbors.

[2] Pixels on the image edge should be consider to have only 5 neighbors, and pixels in image corners should be considered to have only 3 neighbors.

Topics: Achromatic Vision, Gamma, and Visual MTF

Spring 2009 Exam 2: Problem 3 (gamma correction)

Let $T[\cdot]$ be the gamma correction function for $\gamma = 2.2$, and let $T^{-1}[\cdot]$ be its inverse. Furthermore, let $X(m,n)$ be a gray scale image which is linear in energy scaled to the $[0,1]$ range; assume that $T[1] = 1$ and $T[0] = 0$; and also assume that $X(m,n) = 1$ corresponds to white, and $X(m,n) = 0$ corresponds to black.

Then the gamma corrected version of the image is given by

$$\tilde{X}(m,n) = T[X(m,n)]$$

From this data, two different images are formed.

$$\begin{aligned} \tilde{Y}_1(m,n) &= h(m,n) * \tilde{X}(m,n) \\ Y_2(m,n) &= h(m,n) * X(m,n) \end{aligned}$$

where $*$ denotes 2D convolution, and $h(m,n)$ is a low pass filter with an approximate cut-off at $\sqrt{\mu^2 + \nu^2} = \pi/100$. The result, $Y_2(m,n)$, is then gamma corrected to form

$$\tilde{Y}_2(m,n) = T[Y_2(m,n)] \ .$$

For all problems, assume that all displays have a gamma of $\gamma = 2.2$.

a) Assuming that we use a conventional FIR filter implementation, approximately how many multiplies per pixel will it take to implement this filter?

b) Sketch a plot of the gamma correction function $y = T[x]$, for x in the range of $[0,1]$.

c) In general, which image is brighter when displayed, $\tilde{Y}_1(m,n)$ or $\tilde{Y}_2(m,n)$? Why?

d) Assuming that $h(m,n)$ is used to represent the MTF of the human visual system, which of the two images, $\tilde{Y}_1(m,n)$ or $\tilde{Y}_2(m,n)$, would you expect to more accurately match the original image $\tilde{X}(m,n)$ when displayed on a calibrated monitor?

e) Justify your answer to part d).

Spring 2007 Exam 2: Problem 2 (contrast)

Consider an image display system where Y represents the luminance of the output light in units proportional to energy.

When the background luminance is $Y = 10$ a viewer can just notice a spot when it has a luminance of $Y_s = 10.1$, and it has a diameter of 10 degrees in units of visual subtended angle.

a) What is the just noticeable contrast?

b) What is the contrast sensitivity?

c) Assuming that Weber's Law holds true, what luminance must the spot have in order to be just noticeable when the background luminance is $Y = 1$?

d) Select a function $f(Y)$ so that equal quantization steps in the value of $f(Y)$ represent equal changes in contrast. Justify your selection.

Spring 2004 Midterm Exam: Problem 4 (MTF and gamma correction)

Specify a system based on a simple image fidelity model for achromatic images. The systems should:

- Have two inputs consisting of two γ-corrected images, with $\gamma = 2.2$.
- Account for the MTF of the human visual system.
- Account for perceptual sensitivity to contrast.
- Have a single scalar output.

a) Give a block diagram for this system, and specify each block's operation.

b) Explain why each major component is required. When appropriate, give examples of what would go wrong if a component was not used.

c) Give examples of an application where this system might be useful.

Topics: Color matching, additive and subtractive color

Spring 2010 Exam 2: Problem 2 (color and gamma)

Consider the color display that produces the following color (X, Y, Z) when given the input is (r, g, b).

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} (r/255)^\alpha \\ (g/255)^\alpha \\ (b/255)^\alpha \end{bmatrix}$$

a) What is the gamma of the display?

b) What is the white point of the display?

c) What are the color primaries of the display?

d) Can such a display be physically built? Why or why not?

e) You produce two images, one with a checker board pattern of values 0 and 255, and a second with a constant value of $(r, g, b) = (a, a, a)$. The value of a is then adjusted so that the two images match when viewed from a distance. What is the value of gamma for the display in terms of the value a?

f) Imagine that $\alpha = 1$ and the values of (r, g, b) are quantized to 8-bits. Describe the defects you would expect to see in the displayed image.

Spring 2006 Exam 2: Problem 2 (color matching)

Consider the color matching experiment represented by the following diagram.

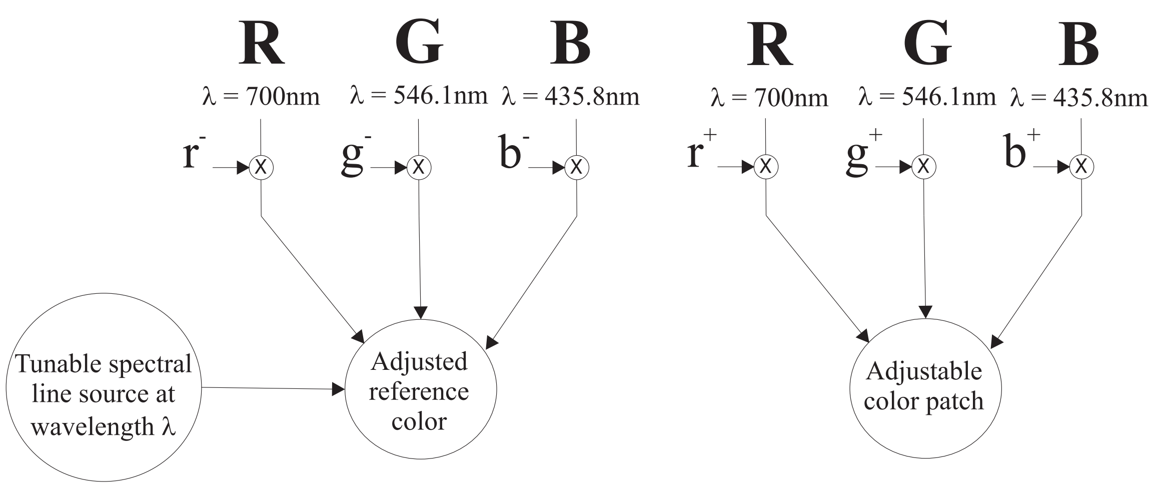

a) Why is it necessary to have both the values of (r^+, g^+, b^+) and (r^-, g^-, b^-) for this matching experiment?

b) For this part, assume that the reference spectral line source to be matched before adjustment is at 480 nm. Indicate the approximate location of the primary colors and the reference line source on the following chromaticity diagram.

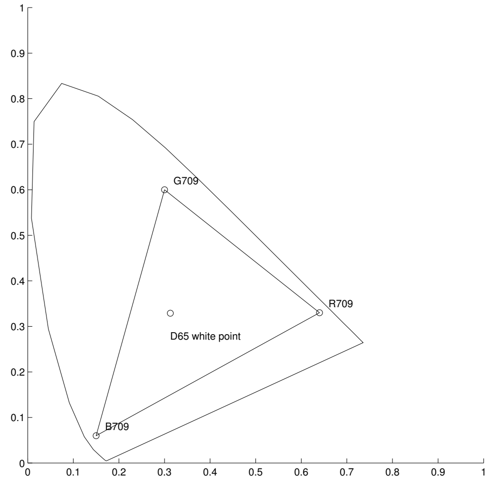

c) For this part, assume that the reference spectral line source to be matched before adjustment is at 480 nm. Which values of (r^+, g^+, b^+) and (r^-, g^-, b^-) will be 0, positive, or negative?

d) Specify how one would measure the color matching functions $r(\lambda)$, $g(\lambda)$, $b(\lambda)$ for the three primaries given in this experiment?

Spring 2003 Exam 2: Problem 3 (subtractive color)

In the following problem, we assume that spectral light measurements are discretized into 31 component vectors ranging from 400 nm to 700 nm in 10 nm steps. Using this assumption, the light reflected from an object has the spectrum

$$I_i = R_i S_i$$

where $1 \leq i \leq 31$ and S_i is the source illumination, R_i is the surface reflectance, and I_i is the reflected light. Further define, x_i, y_i, and z_i as the color matching functions for the X, Y, Z tristimulus values.

For documents printed on a *PurdueJet* printer, it is known that the spectral surface reflectance is given by

$$R = \begin{bmatrix} R_1 \\ \vdots \\ R_{31} \end{bmatrix} = \mathbf{1} - \mathbf{A} \begin{bmatrix} c \\ m \\ y \end{bmatrix}$$

where the columns of the matrix $\mathbf{A}$ are the spectral absorptance's of the cyan, magenta, and yellow inks respectively, and $\mathbf{1}$ is the 31 dimensional column vector $[1, \cdots, 1]^t$.

a) Calculate an equation for the X, Y, Z components of the reflected light in terms of I_i.

b) In general, is it possible for two different surface reflectance functions R' and R'' to have the same X, Y, Z components? Characterize the space of possible spectral differences $\Delta R = R'' - R'$ that will result in no change of the X, Y, Z components.

c) For documents printed on a *PurdueJet* printer, calculate an expression for the vector $[c, m, y]^t$ as a function of the measured value of $[X, Y, Z]^t$ and the known illuminant S. (Hint: You will need to define matrices in terms of the color matching functions and the known illuminant.)

d) For documents printed on a *PurdueJet* printer, calculate an expression for the spectral reflectance vector R as a function of the measured value of $[X, Y, Z]^t$ and the known illuminant S.

Topics: Chromaticity, white point, and quality metrics

Spring 2010 Final: Problem 2 (Lab color transform)

The approximate Lab color space transform is given by

$$\begin{aligned} L &= 100(Y/Y_0)^{1/3} \\ a &= 500\left[(X/X_0)^{1/3} - (Y/Y_0)^{1/3}\right] \\ b &= 200\left[(Y/Y_0)^{1/3} - (Z/Z_0)^{1/3}\right] \end{aligned}$$

a) Qualitatively specify the colors corresponding to the following values of L, a, b: 1) $L = 100$, $a = 0$, $b = 0$; 2) $L = 100$, $a =$ large positive, $b = 0$; 3) $L = 100$, $a = 0$, $b =$ large positive; 4) $L = 100$, $a =$ large negative, $b = 0$; and 5) $L = 100$, $a = 0$, $b =$ large negative;

b) Which component of the L, a, b coordinate system requires the greatest spatial frequency to accurately represent? Why?

c) Is the L, a, b coordinate system suitable for representing images that will be JPEG compressed? Why or why not?

d) Imagine that an image with large energy in high frequencies is viewed from a great distance, and you would like to know the average color that the viewer sees.

Should you low pass filter the L, a, b image to determine the average color? Either justify that this approach is correct, or explain a better approach.

Spring 2010 Final: Problem 1 (chromaticity diagram)

Consider the standard chromaticity diagram below, and assume that you are using a display device with standard 709 r, g, b color primaries.

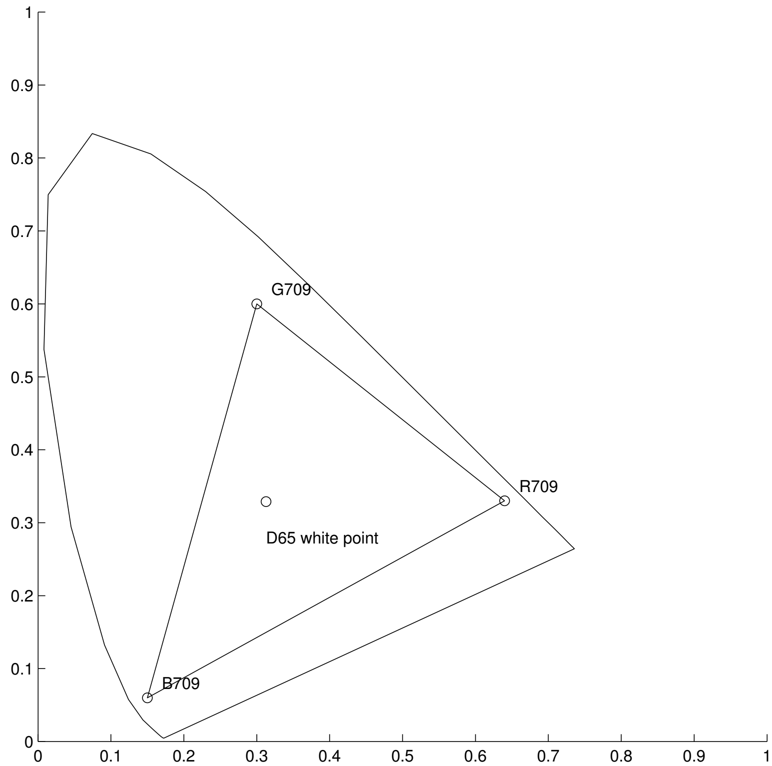

a) Draw a triangle corresponding to all colors with positive values of X, Y, and Z. Label the three primaries for this triangle as "X-primary", "Y-primary", and "Z-primary".

b) Label the region of the chromaticity diagram corresponding to imaginary colors. (Use 45 deg diagonal hash marks to indicate this region of the diagram.)

c) Label ALL real colors with $r < 0$ on the chromaticity diagram. (Use -45 deg diagonal hash marks to indicate this region of the diagram.)

d) Place a point on the diagram corresponding to a highly saturated color that is NOT formed by a single wavelength of light. Label this point with the letter "P".

e) Draw a line on the plot corresponding to all color that can be formed with a combination of the D65 white and R709.

Spring 2009 Exam 2: Problem 2 (chromaticity diagram)

Consider the standard chromaticity diagram below, and for all questions assume that standard 709 r, g, b color primaries are used.

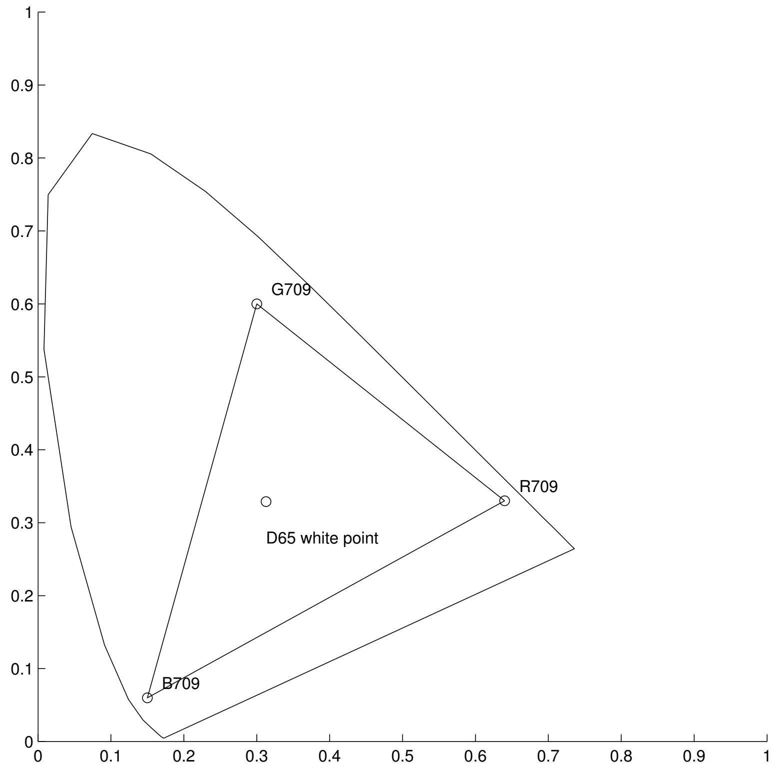

a) Draw the region on the diagram corresponding to $r < 0$, $g > 0$, $b > 0$, and label this region "negative red".

b) Draw a point corresponding to the color primaries for X, Y, Z, and label the three points "X-primary", "Y-primary", and "Z-primary".

c) Draw a triangle corresponding to the gamut of the X, Y, Z color system, when the three tristimulus values are assumed positive.

d) Do all positive values of X, Y, Z correspond to real colors? Why or why not?

e) Do all real colors correspond to positive values of X, Y, Z? Why or why not?

Topics: Color spaces, and perceptual uniformity

Spring 2006 Exam 2: Problem 1 (Lab space)

The approximate *Lab* color space transform is given by

$$\begin{aligned} L &= 100(Y/Y_0)^{1/3} \\ a &= 500\left[(X/X_0)^{1/3} - (Y/Y_0)^{1/3}\right] \\ b &= 200\left[(Y/Y_0)^{1/3} - (Z/Z_0)^{1/3}\right] \end{aligned}$$

a) What are the values of X_0, Y_0, and Z_0 supposed to represent? Typically, what values might be used for these constants? (You don't have to give the specific decimal quantities, just specify a typical choice.)

b) Why are the quantities (X/X_0), (Y/Y_0), and (Z/Z_0) raised to the power 1/3?

c) Why are the quantities $(X/X_0)^{1/3}$ and $(Y/Y_0)^{1/3}$ subtracted for the calculation of a, and the quantities $(Y/Y_0)^{1/3}$ and $(Z/Z_0)^{1/3}$ subtracted for the calculation of b?

d) Specify an application for which the *Lab* color space is well suited. Why?

e) Specify an application for which the *Lab* color space is poorly suited. Why?

Spring 2006 Final: Problem 2 (color transformations)

Consider the color space given by

$$\begin{bmatrix} r \\ g \\ b \end{bmatrix} = M \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

where X, Y, and Z denote the standard CIE color space, and $[r, g, b]$ are the tristimulus values corresponding to physically displayed red (R), green (G), and blue (B) color primaries.

a) Are the entries in the matrix M positive (i.e. ≥ 0), negative (i.e. < 0) or a combination of positive and negative. Justify your answer.

b) Are the entries in the matrix M^{-1} positive (i.e. ≥ 0), negative (i.e. < 0) or a combination of positive and negative. Justify your answer.

c) What does the second column of M^{-1} represent?

d) What does it tell you if the rows of M^{-1} sum to 1?

Topics: Interpolation, decimation, and optimum linear filtering

Spring 2010 Final: Problem 4 (M-estimators and nonlinear filters)

Consider a sequence of N i.i.d. random variables, X_n, each with density

$$p(x|\mu) = \frac{1}{z} \exp\left(-\rho(x-\mu)\right)$$

where μ is a parameter of the distribution, and z is a normalizing constant given by

$$z = \int_{\Re} \exp\left(-\rho(x)\right) dx \ .$$

Then the maximum likelihood estimate of μ is defined as

$$\begin{aligned} \hat{\mu} &= \arg\max_{\mu} \left\{ p(x_1, \cdots, x_N|\mu) \right\} \\ &= \arg\max_{\mu} \left\{ \log p(x_1, \cdots, x_N|\mu) \right\} \end{aligned}$$

where $p(x_1, \cdots, x_N|\mu)$ is the joint density for the sequence of random variables $(X_1, \cdots, X_N)$.

a) Derive an expressions for the joint density, $p(x_1, \cdots, x_N|\mu)$, and $\log p(x_1, \cdots, x_N|\mu)$.

b) Derive a general expression for the maximum likelihood estimate of μ.

c) Calculate the maximum likelihood estimate of μ when $\rho(x) = x^2$.

d) Calculate the maximum likelihood estimate of μ when $\rho(x) = |x|$.

e) What is the advantage of using $\rho(x) = |x|$ rather than $\rho(x) = x^2$?

Spring 2009 Final: Problem 4 (Training for MMSE and LS estimation)

Consider a non-linear prediction problem for which we are trying to predict the value of a scalar Y_n from a vector of observations Z_n. Our assumption is that we can estimate Y_n using the non-linear predictor given by

$$\hat{Y}_n = f(Z_n, \theta)$$

where $\theta \in \Re^p$ is a p dimensional parameter vector that controls the behavior of the nonlinear predictor.

Fortunately, we are given some training data pairs with the form (Y_n, Z_n).[3] The data is partitioned into two sets. The first set, $n \in S_1$, contains $N = |S_1|$ pairs, and is used for training purposes. The second set, $n \in S_2$, contains $M = |S_2|$ pairs, and is used for testing purposes.

[3] Assume that each training data pair is independent, and each pairs has the same distribution.

Using these data, we can define the training MSE as

$$MSE_1(\theta) = \frac{1}{N} \sum_{n \in S_1} \| Y_n - f(Z_n, \theta) \|^2 ,$$

the testing MSE as

$$MSE_2(\theta) = \frac{1}{M} \sum_{n \in S_2} \| Y_n - f(Z_n, \theta) \|^2 ,$$

and the expected MSE as

$$MSE_3(\theta) = E \left[\| Y_n - f(Z_n, \theta) \|^2 \right] .$$

Based on these error measures, we can define the following two estimates for the parameter vector.

$$\hat{\theta} = \arg\min_{\theta} MSE_1(\theta)$$

$$\theta^* = \arg\min_{\theta} MSE_3(\theta)$$

a) Which of the two quantities would you expect to be smaller, $MSE_2(\hat{\theta})$ or $MSE_2(\theta^*)$? Why?

b) What is the disadvantage of using $MSE_2(\theta^*)$?

c) Approximately how large should N be in order for $\hat{\theta}$ to be useful?

d) Sketch the plots of $MSE_1(\hat{\theta})$, $MSE_2(\hat{\theta})$, and $MSE_2(\theta^*)$ as a function of the amount of training data N.

e) Which value would you expect to be smaller, $MSE_1(\hat{\theta})$ or $MSE_2(\hat{\theta})$. Why?

f) If you are reporting results of your experiment, which value should you report, $MSE_1(\hat{\theta})$ or $MSE_2(\hat{\theta})$. Why?

Spring 2009 Final: Problem 3 (MMSE prediction)

Let $Y \in \Re^N$ be a vector containing the pixels in an image window. We can model Y as

$$Y = tS + W$$

where $t \in \Re^N$ is a deterministic column vector of length N, S is scalar valued Gaussian random variable with mean 0 and variance σ^2, and W is a independent Gaussian random vector of correlated noise with distribution $N(0, R_w)$ where R_w is an $N \times N$ positive definite covariance matrix.

Intuitively, Y is composed of a signal tS obscured by noise W. Our objective is to estimate the value of S from the observations Y. To do this, we will form a MMSE linear estimator for S given by

$$\hat{S} = Y^t \theta$$

where $\theta \in \Re^N$ is a vector of coefficients.

Furthermore, define the covariance matrix of Y given by

$$R_y = E\left[YY^t\right] \ ,$$

and the cross-covariance column vector of Y and S given by

$$b = E\left[YS\right] \ .$$

a) Calculate an expression for the MSE given by $E\left[||S-\hat{S}||^2\right]$ in terms of R_y, b, σ^2, and θ.

b) Use the expression from part a) to compute the value of θ that produces the MMSE estimate of S.

c) Calculate R_y in terms of t, σ^2, and R_w.

d) Calculate b in terms of t, σ^2, and R_w.

e) Use the above results to calculate a closed form expression for $\hat{S}$.

Spring 2002 Final: Problem 3 (2D interpolation)

Let the image $y(m,n)$ be formed by applying 2-D interpolation by a factor of $L=2$ to the signal $x(m,n)$ with an interpolation filter of the form

$$\begin{aligned} h(m,n) &= 0.25\delta(m-1,n-1) + 0.5\delta(m,n-1) + 0.25\delta(m+1,n-1) \\ &\quad +0.5\delta(m-1,n) + \delta(m,n) + 0.5\delta(m+1,n) \\ &\quad +0.25\delta(m-1,n+1) + 0.5\delta(m,n+1) + 0.25\delta(m+1,n+1) \end{aligned}$$

a) Use a free boundary condition to compute $y(m,n)$ for the input $x(m,n)$ given by

$$\begin{matrix} 0 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{matrix}$$

b) Compute $H(e^{j\mu}, e^{j\nu})$ the DSFT of the filter $h(m,n)$.

c) Write an expression for $Y(e^{j\mu}, e^{j\nu})$ in terms of $X(e^{j\mu}, e^{j\nu})$ and $H(e^{j\mu}, e^{j\nu})$.

d) What are the advantages and disadvantages of this interpolation method?

Topics: Nonlinear filtering

Spring 2007 Final: Problem 5 (nonlinear filtering)

Consider an signal $Y_n = S_n + W_n$ where S_n is a unknown signal, W_n is i.i.d. Gaussian noise with mean 0 and variance 1.

Your job (should you choose to accept it) is to recover a good estimate of S_n by applying a function to a 5 point window about the location n. So the estimate is given by

$$\hat{S}_n = f(Z_n)$$

where

$$Z_n = [Y_{n-2}, Y_{n-1}, Y_n, Y_{n+1}, Y_{n+2}]^t \ .$$

a) Assuming that $S_n = \mu$, where μ is a constant, then what is a good choice for the function $f(\cdot)$? Justify your answer.

b) Assuming that S_n is a slowly varying function of n, then what is a good choice for the function $f(\cdot)$? Justify your answer.

c) Assuming the S_n has the form

$$S_n = a_n + (10)b_n$$

where a_n is a slowly varying function of n, and b_n is i.i.d. with discrete probability density function $P\{b_n = 1\} = P\{b_n = -1\} = 0.001$ and $P\{b_n = 0\} = 0.998$, then what is a good choice for the function $f(\cdot)$? Justify your answer.

d) Assuming the S_n has the form

$$S_n = a_n + 10 \sum_{k=-\infty}^{\infty} b_k \text{pulse}_{10}(n-k)$$

where a_n is a very slowly varying function of n, and b_n is i. i. d. with discrete probability density function $P\{b_n = 1\} = P\{b_n = -1\} = 0.001$ and $P\{b_n = 0\} = 0.998$, then what is a good choice for the function $f(\cdot)$? [4] Justify your answer.

Spring 2006 Final: Problem 4 (M-estimators)

Consider the set of data $\{x_n\}_{n=0}^{N-1}$ for N odd. We would like to estimate a "central value" using a method known as M-estimation. To do this we compute the following function

$$\hat{\theta} \quad = \quad \arg\min_{\theta} \left\{ \sum_{n=0}^{N-1} \rho(x_n - \theta) \right\}$$

where ρ is a function with the properties that $\rho(\Delta) \geq 0$ and $\rho(-\Delta) = \rho(\Delta)$.

[4] Note that $\text{pulse}_N(n) = u(n) - u(n-N)$.

a) What function $\rho(\Delta)$ will result in the mean as shown below?

$$\text{mean is } \hat{\theta} = \frac{1}{N} \sum_{i=0}^{N-1} x_i$$

b) What function $\rho(\Delta)$ will result in the median?

c) Select a function $\rho(\Delta)$ which usually produces an estimate close to the mean, but limits the influence of a single value of x_i.

Topics: Digital Halftoning

Spring 2010 Final: Problem 3 (quality metrics for halftoned images)

Assume that for achromatic images, you use the following fidelity metric for the human visual system,

$$D = \sum_{m,n} \left(h(m,n) * b(m,n) - h(m,n) * g(m,n) \right)^2$$

where D is a measure of distortion between the linear gray scale image $g(m,n)$ and the binary image $b(m,n)$, and $*$ indicates 2D convolution.

Furthermore, assume that

$$h(m,n) = [\delta(m,n) + (1/2)\left(\delta(m-1,n) + \delta(m+1,n)\right)] * [\delta(m,n) + (1/2)\left(\delta(m,n-1) + \delta(m,n+1)\right)]$$

where $*$ represents 2D convolution.

a) Calculate the DSFT, $H(e^{j\mu}, e^{j\nu})$ of $h(m,n)$, and sketch its shape.

b) What are the values of $H(e^{j0}, e^{j0})$, $H(e^{j\pi}, e^{j0})$, $H(e^{-j\pi}, e^{j0})$, $H(e^{j0}, e^{j\pi})$, and $H(e^{j0}, e^{-j\pi})$.

c) Assuming your objective is to represent the gray scale image $g(m,n)$ by the binary image $b(m,n)$, then is it best for $d(m,n) = b(m,n) - g(m,n)$ to contain mostly high frequencies or low frequencies? Why?

d) If $g(m,n) = 1/2$, then determine all binary patterns $b(m,n)$ (i.e. a pattern of 1's and 0's) that best matches $g(m,n)$.

e) How could the distortion measure, D, be improved to better account for contrast?

Spring 2009 Final: Problem 5 (error diffusion)

Consider the 1-D error diffusion algorithm specified by the equations

$$\begin{aligned} b_n &= Q(y_n) \\ e_n &= y_n - b_n \\ y_n &= x_n + e_{n-1} \end{aligned}$$

where x_n is the input, b_n is the output, and $Q(\cdot)$ is a binary quantizer with the form

$$Q(y) = \begin{cases} 1 & \text{if } y > 0.5 \\ 0 & \text{if } y \le 0.5 \end{cases} .$$

where we assume that $e_0 = 0$ and the algorithm is run for $n \ge 1$.

Furthermore, define $d_n = x_n - b_n$.

a) Draw a flow diagram for this algorithm. Make sure to label all the signals in the flow diagram using the notation defined above.

b) Calculate b_n for $n = 1$ to 10 when $x_n = 0.25$ and $e_0 = 0$.

c) Calculate an expression for d_n in terms of the quantization error e_n.

d) Calculate an expression for $\sum_{n=1}^{N} d_n$ in terms of the quantization error e_n.

e) What does the result of part d) above tell you about the output of error diffusion?

Spring 2004 Final: Problem 2 (halftoning/power spectrum)

Let $X(m, n)$ be an achromatic image taking continous values in the interval $[0, 1]$, and let $T(m, n)$ be a 2-D random field of i.i.d. random variables which are uniformly distributed on the interval $[0, 1]$. Let the halftoned version of $X(m, n)$ be given by

$$Y(m, n) = \begin{cases} 1 & \text{if } X(m, n) \geq T(m, n) \\ 0 & \text{if } X(m, n) < T(m, n) \end{cases}$$

a) Is $Y(m, n)$ a stationary random process?

b) Calculate $\mu(m, n) = E[Y(m, n)]$

c) Calculate $E[D(m, n)D(m + k, n + l)]$ for $D(m, n) = Y(m, n) - \mu(m, n)$.

d) Is $D(m, n)$ a stationary random process?

e) For the special case of $X(m, n) = g$, compute the power spectral density $S(e^{j\mu}, e^{j\nu})$ of $D(m, n)$.

f) Is $Y(m, n)$ a good quality halftone of $X(m, n)$. Justify your answer.

Topics: Entropy and lossless image coding

Spring 2008 Final: Problem 2 (entropy coding)

Let X_n be a discrete-time random process with i.i.d. samples, and distribution given by $P\{X_n = k\} = p_k$ where

$$(p_0, p_1, p_2, p_3, p_4, p_5, p_6, p_7) = \left(\frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{32}, \frac{1}{4}, \frac{1}{8}, \frac{1}{8}, \frac{1}{4}\right)$$

a) What is the value of p_8? Why?

b) Calculate the entropy $H(X_n)$ in bits.

c) Draw the Huffman tree and determine the binary Huffman code for each possible symbol.

d) Calculate the expected code length per symbol.

e) Are there better codes for X_n? If so, what are they? If not, why not?

Spring 2005 Final: Problem 1 (lossless image coding)

Consider a lossless predictive coder which predicts the pixel $X_{s_1,s_2} = k$ from the two pixels $X_{s_1,s_2-1} = i$ and $X_{s_1-1,s_2} = j$. In order to design the predictor, you first measure the histogram for the values of i, j, k from some sample images. This results in the following measurements.

i	j	k	$h(i,j,k)$
0	0	0	30
0	0	1	2
0	1	0	4
0	1	1	12
1	0	0	12
1	0	1	4
1	1	0	2
1	1	1	30

a) Use the values of $h(i, j, k)$ to calculate $\hat{p}(k|i, j)$, an estimate of

$$p(k|i,j) = P\left\{X_{s_1,s_2} = k | X_{s_1,s_2-1} = i,\ X_{s_1-1,s_2} = j\right\}$$

and use them to fill in the table below.

i	j	k	$\hat{p}(k\|i,j)$
0	0	0	
0	0	1	
0	1	0	
0	1	1	
1	0	0	
1	0	1	
1	1	0	
1	1	1	

b) For each value of i, and j, compute a binary valued estimate of X_{s_1,s_2} and use it to fill in the table below.

i	j	$\hat{X}_{s_1,s_2}$
0	0	
0	1	
1	0	
1	1	

c) Draw a block diagram for the lossless predictive coder. The block diagram should include an entropy coder.

d) Assuming the prediction errors are independent (but not identically distributed), calculate an expression for the theoretically achievable bit rate for the lossless predictive encoder. Justify your answer. (Hint: Use the fact that $\log_2(3) = 1.585$, $\log_2(7) = 2.807$, and $\log_2(15) = 3.907$.)

Spring 2001 Final: Problem 1 (entropy coding)

Let X_n be a discrete random variable which takes values on the set $\{0, 1, \cdots, 5\}$, and let

$$P\{X_n = k\} = p_k$$

where

$$(p_0, p_1, p_2, p_3, p_4, p_5) = (0.1, 0.04, 0.2, 0.06, 0.35, 0.25)$$

a) Draw and fully label the binary tree used to form a Huffman code for X_n.

b) Write out the Huffman codes for the six symbols $0, 1, \cdots, 5$

c) Compute the expected code length for your Huffman code.

Topics: Lossy image source coding

Spring 2008 Final: Problem 4 (rate-distortion))

Consider a discrete-time random process X_n with i.i.d. samples that are Gaussian with mean 0 and variance $\sigma^2 > 0$.

The rate distortion relation for this source is then given by

$$\begin{aligned} R(\Delta) &= \max\left\{\frac{1}{2}\log_2\left(\frac{\sigma^2}{\Delta^2}\right), 0\right\} \\ D(\Delta) &= \min\left\{\sigma^2, \Delta^2\right\} \end{aligned}$$

a) Plot the minimum possible rate (y-axis) versus distortion (x-axis) required to code this source when $\sigma^2 = 1$.

b) If we require that the distortion $D \leq \sigma^2$, then what is the minimum (lower bound) on the number of bits per sample that is required to transmit this signal?

c) If we require that the distortion $D \leq \frac{\sigma^2}{(32)^2}$, then what is the minimum (lower bound) on the number of bits per sample that is required to transmit this signal?

d) How many bits per sample are required in order to achieve zero distortion?

e) Describe how you would design a lossy coder for this signal assuming that your objective is to achieve a bit rate of approximately 8 bits per sample.