Midterm Exam

Assigned: 10/10/02
Due: 10/22/02 at 2:30pm

Problem 1 Color Constancy (Matlab)

The Matlab data file `constancy.mat` contains several different spectral response vectors which will be needed for this problem. The vectors s1, s2, s3 specify sensitivity as a function of wavelength for three different sensor types. For all spectra used in this problem, the response curves are sampled at 5 nm intervals for wavelengths from 400 through 700 nm. These sampling wavelengths are stored in the lambda vector.

Assume that surface reflectance spectra are well fit with a 3 dimensional linear model. The basis functions for those spectra are given by vectors b1, b2, b3 (so that every physical surface reflectance spectrum can be written as some linear combination of b1, b2, and b3). These basis functions were obtained via principal components analysis of empirical observations, so that b1 models the most significant variations, b2 the second most, and b3 the least significant. Assume that illumination spectra are also well fit by a (different) PCA-based 3 dimensional linear model, with basis vectors i1, i2, i3.

We examine readings taken from 20 color patches using 3 different cameras, each containing a different number of sensors. Whenever asked to determine the illumination or surface spectra, you should specify these in terms of the appropriate spectral basis functions. You should always make as few assumptions as possible. If you need to use the gray world assumption to solve the problem, do so (assume that gray corresponds to all sensors giving a response of 0.5). If additional constraints are needed to solve the problem, specify a sufficient set of additional constraints.

a) **3-sensor camera.** The camera readings from sensors s1, s2, s3 for each of 20 different surface patches, averaged over the pixels of each patch, are given in the matrix P (each row is a different patch, each column a different sensor). Find the illumination and surface spectral basis function coefficients.
b) **2-sensor camera.** Now assume that only the first two sensors are available (i.e. use only the first two columns of P). Find the illumination and surface spectral basis function coefficients. *Hint:* Consider how the numbers of equations and unknowns change as the number of spectral basis functions is varied.
c) **4-sensor camera.** Suppose that a fourth sensor is available, with a spectral response s4 that is linearly independent of the first three sensors. How many constraints do the 20 color patch observations impose on the illumination and reflectance basis function coefficients? What is the minimum number of noise free observations needed to calculate the surface and illuminant spectra? Is there an analytic solution? If so, give it.

Problem 2 CIE Color Coordinates (Matlab)

The Matlab data file `CIE.mat` contains the spectra which will be needed for this problem. The vectors cx, cy, and cz give the CIE color matching functions for the X, Y, and Z color coordinates, respectively (these are also plotted in Figure 6.7 of Forsyth &

Ponce). As in problem 1, these functions are sampled at 5 nm intervals for wavelengths from 400 through 700 nm.

a) What are the CIE X, Y, and Z coordinates of the light corresponding to the power spectrum in the vector sA? What are the normalized CIE coordinates, x and y?
b) Suppose you have a test color specified as a linear combination of primaries at 450, 550, and 650 nm, given by $(a, b, c)^T$, respectively. What linear combination of light sources at 470, 525, and 600 nm is need to match the test color?
c) We want to specify one possible spectral power distribution of the non-physical primary lights of the CIE color coordinate system, corresponding to the color matching functions cx, cy, cz. Suppose you insist that the spectral power of the CIE primaries be zero for all wavelengths except at 450, 550, & 650 nm. Specify the linear combination of lights at 450, 550, & 650 nm that corresponds to the primary lights of the CIE color coordinate system.

Problem 3 Small Camera Motions

Exercise 10.2 from *Computer Vision: A Modern Approach* (Forsyth and Ponce)

Problem 4 Correlation Stereo Matching

Exercise 11.5 from *Computer Vision: A Modern Approach* (Forsyth and Ponce)

Problem 5 Recursive Computation of the Correlation Function

a) Recursive filtering techniques are often used to reduce the computational complexity of a repeated operation such as filtering. If an image filter is applied to each location in an image, a (horizontally) recursive formulation of the filtering operation expresses the result at location *(x+1,y)* in terms of the previously computed result at location *(x,y)*. Consider a box convolution filter, which has coefficients equal to one inside a rectangular window, and zero elsewhere:

$$B(x, y, w, h) = \sum_{j=0}^{h-1} \sum_{i=0}^{w-1} I(x+i, y+j)$$

Design a new box filtering method which is applied recursively along particular rows of the image (incrementing x) and efficiently reuses the value computed at the previous location.

b) Solve exercise 11.7 from *Computer Vision: A Modern Approach* (Forsyth and Ponce). Note that in the equations for this problem, $\bar{I}$ is the average value within the correlation window (*not* the average over the whole image), and $\bar{w}$ is a constant vector with each component equal to $\bar{I}$. Consider only recursion in one dimension.

c) An integral image allows for even faster computation of large box filters. The integral image *J* at location *(x,y)* is defined to be the sum of the pixel values above and to the left of *(x,y)*, inclusive:

$$J(x, y) = \sum_{j=0}^{y} \sum_{i=0}^{x} I(i, j)$$

Show how you can recursively compute an integral image.

d) Show that, given an integral image computed from an input image, the value of a box filter anywhere on the original image can be computed using four references to the integral image.

$$B_J(x, y, w, h) = aJ(?,?) + bJ(?,?) + cJ(?,?) + dJ(?,?)$$

Specify what the ?'s and the scalars *a, b, c, d* should be to make this formula correct.

e) In what cases can the integral image representation help reduce the computation required for normalized correlation beyond what you concluded in part (b)?

Problem 6 Bayesian Inference: The Skyrise Lights Problem

You see the word "HI" spelled out in lights on the side of a building, and you wonder whether this came about by chance:

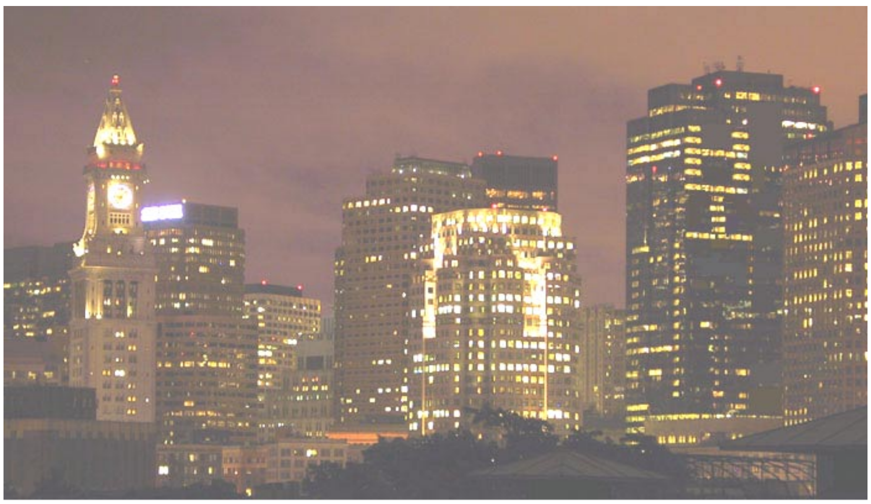

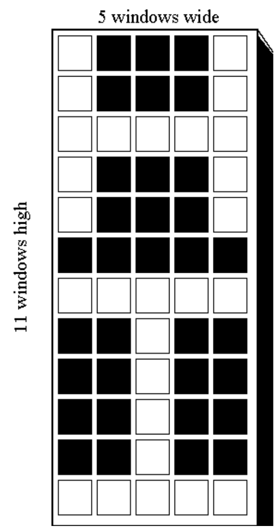

Throughout this problem, we consider only the 11x5 grid of windows shown in the figure at the right. You know of 3 different possible processes for setting the lights in the tower:

(i) The setting of the lights in each room is independent and random, with the probability of a light being left on equal to 0.4.

(ii) One night every month, the ballroom dance club meets in the tower. The club requires an entire floor of five rooms to be used (and lit) for each instruction level. There are three dance instruction levels each month, each meeting on a different, randomly selected floor. The rest of the rooms in the building are lit or not, according to the random process of (i).

(iii) You believe that about one night every 10 years (not on a dance club night), someone will orchestrate a hack to spell out some two-letter word on the building, setting the lights in each room of the tower by hand. Assume there are 10 different possible two-letter words to be spelled, all equally probable. One of these words is "HI".

Show your work when answering the following questions:

a) Assume that every night's lighting not caused by processes (ii) and (iii) is caused by the completely random process, (i). Based on the number of days over 10 years (ignoring leap years) that you expect each process to be in effect, what is the prior probability for each of the three possible lighting hypotheses?

b) What is the likelihood (in the Bayesian usage) of the observed data, for each of the 3 different hypotheses?
c) What is the posterior probability of the observed data, for each of the 3 different hypotheses? What is the most probable inference, given the observation of "HI" spelled out in lights?

Problem 7 Surface Normals to Surface Shape (Matlab)

For the photometric stereo problem, and for the shape-from-shading algorithm discussed in class, the output shape is represented as a set of surface normals at each x and y position (or derivatives of surface height, z, with respect to x and y). Often one wants to recover the 3-d shape (surface height), not just the surface normals or slopes.

Suppose we have estimates of the quantities p and q, where

$$\frac{d}{dx} z(y,x) = p(y,x) \qquad\qquad \frac{d}{dy} z(y,x) = q(y,x) \tag{1}$$

and we want to compute z, the surface height, as a function of position (y,x). Here, we have specified coordinates as (y,x) to match Matlab's array indexing convention. Using finite differences to represent derivatives, we can rewrite eq. (1) in a matrix form,

$$Az = b \tag{2}$$

where the column vector z is a rasterized version of the surface height matrix, $z(y,x)$. For example, the finite difference approximation to the x derivative $p(y,x)$ would be $(z(y,x+1) - z(y,x))$.

a) For a 4x4 image, write the elements of the matrix A, and specify what b is. For filter outputs at the image boundary, you can assume that all surface values are zero outside of the 4x4 image.
b) For overdetermined systems such as eq. (2) (where A has more rows than columns), the least squares solution is $z = inv(A' * A) * A' * b$. If $A'*A$ is not of full rank, the pseudo-inverse (*pinv(A)*) can be used. Construct a test image with pixels $z(2,2) = 1$, and $z = 0$ everywhere else. Then, compute the derivatives, $p(x,y)$ and $q(x,y)$, and show that the least squares solution recovers a reasonable estimate of the surface from the p and q data. Explain why the surface estimate is a reasonable one.
c) Using the *randn* function, add zero-mean Gaussian measurement noise, of standard deviation 0.3, to the p and q values from part (b) and reestimate the surface height. Explain why it is or isn't a reasonable surface estimate.
d) Now add an additional constraint term to the objective function: that the surface slopes are smooth. Use a quadratic constraint as discussed in class for the shape from shading problem, and include a regularization parameter, λ. Specify the discrete approximation to this regularized objective for the 4x4 image. Compute the least squares solution to the regularized problem for the noisy data from part (c). Compare the results obtained for three different regularization parameters λ = 0.1, 1, 10.