CSE 803 Class,

Here is some more info about the final.

The final is on Monday, Dec. 13, 12:45 - 2:45pm in our same classroom.

Coverage is comprehensive -- the entire course -- with about 2/5 coverage of Chapters 1-7 and 3/5 on the rest of the text. More advice below. Format will be mixed: multiple choice for facts and terms; short discussion of concepts/phenomena; a design problem; a set of mathematical models for various important phenomena/methods.

I would like you to have a good "feel" for CV overall, of course. This means, know some methods, concepts, facts, and phenomena -- and be able to design systems for machine vision. This is a lot! There will be choices as in the midterm.

Prepare some thoughts on the design questions cited below.

Prepare some thoughts about the various units studied, such as stereo, shape-from-shading, convolution, etc. Consider this structure in the case of S-f-X. (Shape-from-?)

> There is a physical phenomena -- what is it?
> (Image shading is the result of illumination interacting with surface material and shape.)
> There is an information processing view -- what is it?
> (We input an intensity image and light source position; we output the surface normals of the object.)
> There is a simple mathematical model -- what is it?
> (intensity = a cos (theta) ...)
> There are methods to compute the output. (1) propagation of normals from object limbs together with above formula, (2) photometric stereo, which is ...

Chapter coverage:

If you understand some major terminal chapters, you should be OK on the prerequisite chapters.

Chs 1 and 2 are motivational with easy concepts and details;
the notions of quantization and resolution are most important.
Ch 3 has a few major concepts: CC and feature computation, mophology
(we lightly covered morphology).
Ch 4 has a few major PR ideas (ANNs not covered)

Ch 5 is quite long, but has only a few major ideas: masks and convolution are extremely important, vector space ideas, Fourier transorm.
Ch 6,7 a few major color ideas and texture ideas
Ch 8 how to use everything previously learned to index to images
Ch 9 a few major motion ideas: recall Project 3
Ch 10 the major notion of image segmentation and some ideas on how to do it and how to represent the results; line and curve extraction; region extraction
Ch 11 several different matching methods
Ch 12 several phenomena supporting 3D perception
Ch 13 several concepts and methods for quantitatively sensing 3D.
Ch 14 several 3D modeling methods and some important concepts of 3D modeling
Ch 16 Iris system.

FINAL EXAM DESIGN QUESTIONS:

At least two of the following questions will be asked, verbatim, on the final exam. You may discuss these problems with others, however, you cannot bring any written material into the final exam -- all must be in your own memory. You may ask questions of the instructor for clarification of issues; however, do not look for an oracle to comment on your rough drafts. If you can handle these questions, then you could work in the field of machine vision.

1) The problem is to make a good 3-class decision to classify an image for image-based retrieval. A the top of the hierarchy, the decision to be made is
   (a) image of nature only (trees, shrubs, grass, rocks, water, sky, etc.
   (b) image of urban scene (buildings, roads, vehicles, etc., and possible a small amount of (a) stuff
   (c) mixed urban and natural (a significant aount of both (a) and (b)

   How would you process the image to make the above decision? What features would be computed and how? What process would make the decision based on the features? (If you do not want to use a feature-based approach, then describe an alternate approach to making this decision.)

2) The problem is to decide how many coins of each type (penny, nickle, dime, quarter) are on a table in front of a color camera. Consider the use of (a) connected components analysis (versus) (b) circular Hough transform to solve this problem. (i) Carefully, describe the steps necessary for implementing each of these processes. How do we go from the input image to the count of each kind of coin? (ii) Compare the relative advantages and disadvantages of these two approaches. (iii) Assuming it is your job to implement one of these approaches, which would you choose and why?

3) The problem is to measure the height h (and width w) of a building from an aircraft. We want to estimate both the measurement and the expected error of the measurement from

two different methods. The figure below illustrates the sensing situation. The aircraft flies at high altitude and takes two images of a flat section of earth at times t1 and t2. Each image has 1000 x 1000 pixels and the nominal resolution on the ground is 1m per pixel. Make the following assumptions. Altitude a=15,000 is known precisely as is the sun angle theta = 45 degrees. The distance between the focal points of image I1 and image I2 is b. Focal length is f=300mm. The orientation of the plane is identical at times t1 and t2 and the optical axis points perpendicularly downward. (You may make other assumptions if you find it necessary: be sure to state them.)

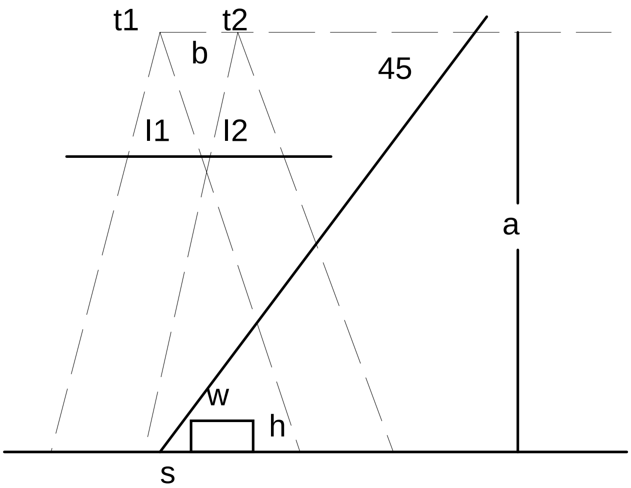

(a) Describe how building height h is computed using stereo from the images I1 and I2.
(b) Discuss how the image of the building will be detected in the image and the expected error in its detection. What is the error in the measurement h as a function of the error in the image of the building? Same question for measurement of w.
(c) It is known that the sun angle is precisely 45 degrees. H can be computed by measuring shadow s in the image. Discuss how this can be done and discuss the error possible in h as a result.
(d) Decide which measurement technique is better and why?

4) The problem is to decide which blocks of which streets need to be plowed to clear snow.

Cities and townships that do snow plowing will contract with the local airlines to use cameras on the planes to provide aerial images of the local area as they fly in and out of Lansing Airport. An image processing system is built that takes these images as input and outputs a set of street blocks to be plowed. (A “block” is a section between two “street corners”.) Digital maps of the area are also available to the system.

(a) Describe how a digital image can be put into correspondence with the corresponding map[s].
(b) Describe how sections/blocks of streets can be found in an image.
(c) Describe how it can be determined that a block needs to be plowed or not plowed.