Practice Midterm

CMSC 426

The midterm will cover material up to and including background subtraction. Here are some topics it will be helpful to master for the midterm:

- Histograms
  - What is a histogram? What is a cumulative distribution function (CDF)? What is histogram equalization and histogram specification? How do you compare histograms using SSD or Chi-squared comparisons? How do you represent a histogram using uniform discretization of a space?
- Correlation
  - How do you perform correlation in 1D or 2D? What is a box filter? How do you create a Gaussian filter? How do you create filters to compute derivatives? Fourier transforms (only for challenge problems).
- Sampling, aliasing and multiscale
- Edge detection
  - Image gradients
  - Non-maximum suppression
- Intelligent Scissors
  - Understand how gradients are computed and used. How shortest path can give good boundaries. How one can incorporate smoothness into IS.
- Human perceptual grouping
  - Basic gestalt grouping cues. How would different theories of grouping explain some of these?
- K-means clustering
  - Understand algorithm. Understand how each iteration reduces the sum of square distances from points to centers.
  - Understand how K-means provides a new way of representing histograms.
- Color
  - RGB and HSV representations. Basics of how the eye senses color.
- Texture
  - Texture represented as histogram of filter outputs. Texture as a Markov process.
- Background subtraction
  - Kernel density estimation.

Practice: The goal of this is to give you samples of the sorts of questions and topics that will come up in the midterm. Some of these questions may be a bit more involved or more vague than those that I would ask in a real midterm. It is very likely that there will be at least one question on the midterm that is quite similar to a practice question.

1. Consider the following two images:

1	1	2	1	2
2	1	2	2	2
3	2	3	3	2
3	3	3	4	3
3	4	3	4	5

5	5	5	5	5
4	5	4	4	4
3	4	3	3	2
2	3	2	2	2
1	2	1	1	1

Compute their histograms and CDFs. Compute the distance between them using SSD and the Chi-squared distance.

2. Give the result of correlating image I with the correlation filter F. Do not include boundary pixels for which the kernel does not completely overlap the image.

I =

3	4	6	9
3	3	4	6
2	3	3	4
2	1	3	3

F =

0	.1	0
.1	.6	.1
0	.1	0

3. Give a 5x5 correlation kernel that will simultaneously smooth the image by averaging with a box filter and take the first derivative of the image in the *x* direction. That is, applying this filter should be equivalent to averaging with a 3x3 box filter and then taking a first derivative.

4. Consider the following, small image:

4	6	9	12	15
3	4	6	9	12

3	3	**4**	6	9
2	3	3	4	6
2	1	3	3	4

a. Compute the magnitude and the direction of the gradient for the central point (**4**).
b. Would the Canny edge detector consider this point to be an edge? Explain why or why not in detail, describing what conditions are needed to decide this.

5. **Texture:** For each pair of the following 1D textures, indicate whether we can distinguish between them using 1) the histogram of the image; 2) a Markov model in which each pixel depends on its immediate neighbors; 3) a Markov model in which each pixel depends on its two immediate neighbors.

A: 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1
B: 1 1 1 0 0 0 0 0 0 1 1 1 1 1 1 0 0 0
C: 0 1 1 0 0 0 1 1 0 1 1 1 0 1 1 0 1 1

6. **Challenge Problem:** Consider an infinite, continuous image *I* where the image is described by the equation: *I(x,y) = x*x + y*y*. This image is convolved with a box filter of width 2.

   a) Use the definition of convolution to write an equation describing the resulting image as the result of taking some integrals.

   b) Solve this equation and write an expression for the resulting image as a function of just *x* and *y*.

For partial credit, solve this problem in 1D, with the image *I(x) = x*x*.

7. Show the result of convolving the kernel, *k*, with the image, *I*.

1/4	1/2	1/4

k

0	0	...	0	1	2	...	10	0	0	...

I

8. Give an example that shows that correlation is not associative

9. Background Subtraction: Suppose you have a video camera set up in front of a concrete wall, taking 30 images per second, and you want to detect when people walk by. The wall is always a uniform grey with an intensity of 128. However,

there is a light snow. Each snow flake appears in a single pixel and turns it to an intensity of 255. About one pixel in 100 has a snow flake in it.

a. Suppose someone walks every few minutes, and we use the background subtraction algorithm described in class, which uses kernel density estimation, and treats every pixel independently. Explain when this algorithm will make errors, and how these errors might depend on the parameters of the system.
b. Suppose the scene grows crowded, so that a new person enters the image every second. What new problems will you have?

10. Construct a 1D filter for a Gaussian with a sigma of 1.