cG89.2223 Weekly Questions (Maloney)

1. David is windsurfing parallel to the beach. Suppose that he has two depth cues to the distance to the beach. Both are distributed as independent Gaussian random variables.

a. What does it mean to say they are unbiased?
b. Assume the cues are unbiased. The variance of the first is 4, that of the second is 1. What weights would an optimal observer minimizing variance give to the two cues? What would the variance of his resulting estimate be?
c. Suppose that David just averages the unbiased two cues (giving them equal weight). What is the variance of his estimate? How much worse is it than the variance of the optimal estimator?
d. Compared to the "average cue" in 1c would David be better off in just using the lower variance cue and discarding the other?

2. Michael is combining cues to estimate the location of a small invisible target circle of radius 1. He marks his estimate with a small dot. If his estimate is in the target circle, he gets a prize. Windsurfing lessons!
The center of the target circle is at (x,y). Michael wants an estimate as close to (x,y) as possible.
He has two estimates of the center of the target, two bivariate Gaussian cues (X1,Y1) and (X2,Y2).
X1 is Gaussian with mean x and variance 1, Y1 is Gaussian with mean y and variance 1.
X2 is Gaussian with mean x and variance 4, Y2 is Gaussian with mean y and variance 4.
All the variables are independent.

a. If Michael uses the more accurate cue (X1,Y1) what is his probability of hitting the target? You can compute this by Monte Carlo (at least 1000000 trials please). Please include your code (MATLAB, LISP, PYTHON, C++, R, whatever).

b. Suppose he combines X1, X2 and Y1,Y2 by forming a weighted average

$$X = w * X1 + (1 - w) * X2$$
$$Y = v * Y1 + (1 - v) * Y2$$

into an estimate X and Y1, Y2 into an estimate Y of y. What should w and v be? What is the probability of hitting the target if he combines the cues to minimize the variance and estimating X and Y?

3. Marisa is looking for Denis in a very large one-dimensional shopping mall. Location is specified by a coordinate X. Marisa knows that, all else being equal, Denis prefers to be near the center of the shopping mall at location 50. He has a prior Gaussian distribution centered on 50 with variance 64. The only clue Marisa has is a coffee cup of a brand that only Denis drinks that she finds at location X=30. The coffee cup is cold and Denis has

wandered off. Based on the location of the coffee cup, the likelihood function of his location is a Gaussian distribution with mean X=30 and variance 100.

a. Explain how you would frame this problem as a problem in Bayesian estimation, using appropriate terminology. What is Denis' posterior distribution? Draw his prior, likelihood and posterior distributions on a single plot.
b. The coffee cup was not that cold after all. Denis' likelihood function has mean X=30 but with a smaller variance 36. Redo part a. Describe what happened to the posterior distribution. Has it moved? Does the change make sense?
(the Stocker & Simoncelli paper can help).

4. Roozbeh wants to try Michael's task in problem 2. But his potential prize is $1000! Moreover, the people running the contest have decided that Roozbeh can have as many cues as he likes! The cues are independent, identically distributed with variance in both the x and y direction of 10. The means of the X values are always x, those of the Y values always y. Unfortunately, the cues cost $10 each. How many cues should Roozbeh purchase? If he purchases none, he will likely miss the target. If he purchases 101 then he has purchased $1010 in cues to have some chance of winning $1000. Bad idea! What is the number of cues that maximizes his expected gain?
The simplest approach would be a Monte Carlo simulation of the task repeated many times. Let us know if your programming skills need a bit of polish to do this problem.