Midterm CMSC828J 2012

Assigned, 3/26/12

Due, 4/4/12 (Turn in to Joao Soares AV Williams 4420)

I made a few clarifications on 3/27/12, shown in red.

1. Let G_σ be a 1D Gaussian filter with s standard deviation of σ. Let u(t) = G_σ*cos(t), that is, the cosine function filtered with the Gaussian. If u(0) = .9, what is the value of u(π/8), u(π/4), u(π/2)?
2. Suppose u(t) = sin(t), for 0<= t <=2π. We perform one iteration of Perona-Malik diffusion on this signal. For simplicity, you may assume that σ = 0, so that there is no smoothing before computing derivatives. You may also ignore what happens at the boundary (t = 0, 2π), or treat this in any way you like (ie., assume the signal is periodic). Note also that this problem concerns a 1D signal. The 1D version of Perona-Malik is the obvious simplification of the 2D version.
    a. Give an analytic expression for the method noise (see Baudes et al.). Plot the noise.
    b. Using the results from (1), how would you expect the result to change for non-zero values of σ.
3. In Baudes et al., we see that non-local means tends to produce method noise that looks like white noise. If we apply non-local means to u(t) = sin(t), will the method noise be like white noise? That is, to get at this, let's consider the method noise at the points t and t + δ, for all values of t. This gives us method noise at a bunch of paired locations. Will this method noise be uncorrelated as δ goes to 0? Provide a proof (or at least a very convincing argument) that your answer is correct.
4. Consider the "map" below, showing the six "regions", A,B,C,D,E, and F. Suppose a rabbit starts in A, and every second, jumps from its current region to one of the neighboring regions randomly, with all neighboring regions being equally likely as the next destination. After a long time has passed, what is the probability distribution of the rabbit's location?

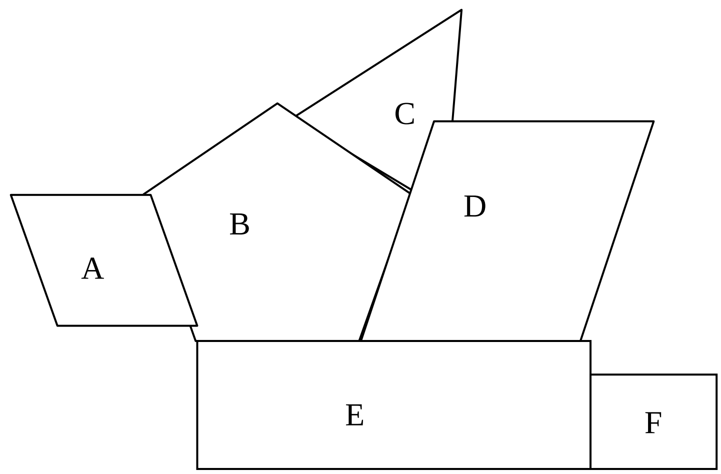

5. Markov Random Fields. Suppose we have an image in which every pixel should be labeled “house”, “boat”, “sea” or “sky”. These classes have the following properties:
   a. In the absence of any other information, all labels are equally likely to apply to a pixel.
   b. A house pixel is red 70% of the time and brown 30% of the time.
   c. A boat pixel is red 50% of the time and brown 50% of the time.
   d. Sea and sky pixels are always blue.
   e. A house pixel is next to another house pixel 95% of the time, next to sky 4% of the time and next to sea 1% of the time.
   f. A boat pixel is next to another boat pixel 95% of the time, next to sky 1% of the time, and next to sea 4% of the time.
   g. Sea and sky pixels are never next to each other.
   h. (Technically, these probabilities may lead to some labelings having 0 probability; I really should make some of the things that have 0 probability have ε probability, for some tiny ε, but this would complicate things. This issue can be safely ignored for this problem).

   Use this information to create a Gibbs Random Field. Describe all potentials fully.

   Suppose that we optimize the labeling of an image using a graph cuts algorithm. Is it possible that we will get caught in a local optimum that is not the globally optimal labeling? Give a proof (or convincing argument) that this cannot happen, or an example of an image and initial labeling that will lead us to get stuck in such a local optima (along with an explanation of why we will be stuck).

6. Suppose we have an NxN image. We create a graph for normalized cut in which each pixel is connected to its four neighbors with an edge weight of 1. Suppose we divide the graph into two parts. One part is an AxB rectangle in which ***none of the pixels are on the boundary of the image***.
   a. What is the normalized cut cost (Ncut) of this separation of the graph?
   b. Suppose we wanted to find the rectangle, as above, that optimizes the Ncut cost. Is this cut optimized by such a rectangle, or is there some other partition of the graph that would optimize this cost? If the answer is yes, how would you modify the graph so that the rectangle that we are looking for is produced by the partition of the new graph that optimizes the Ncut cost?
   c. Prove that of all such rectangles, the one that optimizes the Ncut cost is a square. One thing that can make this problem a little tricky is that A and B really need to be integers. However, if you like, you can formulate an appropriate continuous version of the problem and provide a solution for that version.