Homework 2

Prepare a web page containing answers to the following questions, and email the URL to me by April 5, 2008. Please use the term "homework" in the subject line of your email, and give me the names of all participants. You may work in groups of up to three.

Written items

Prepare an answer to the first question, and at least three of the others.

1. We wish to use EM to estimate a mixture of symmetric Gaussians. Assume that there are n Gaussians in the mixture. Write $\mathbf{x}_i$ for the i'th data point, $\delta_{ij} = 1$ if the i'th data point comes from the j'th Gaussian and zero otherwise, μ_j for the mean of the j'th Gaussian, Σ for the covariance (which is the same for each Gaussian, and is known), and π_j for the weight of the j'th component. Consider the expression:

   $$G(\mu, \delta) = \sum_{ij} \delta_{ij} \left(-(1/2)(\mathbf{x}_i - \mu_j)^T \Sigma^{-1} (\mathbf{x}_i - \mu_j) + \log \pi_j\right) + \sum_{ij} \delta_{ij} \log \delta_{ij}$$

   Show that the EM update equations are equivalent to maximizing this expression, subject to the constraints $\sum_j \pi_j = 1$ and $\sum_j \delta_{ij} = 1$ by repeatedly (a) fixing δ and maximizing G with respect to μ, π subject to relevant constraints and (b) fixing μ, π and maximizing G with respect to δ subject to relevant constraints.

2. Write $\mathcal{E}_0$ for an image that consists of all zeros, with a single one at the center. Show that convolving this image with the kernel

   $$H_{ij} = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{((i-k-1)^2 + (j-k-1)^2)}{2\sigma^2}\right)$$

   (which is a discretised Gaussian) yields a circularly symmetric fuzzy blob.

3. A separable 2D filter kernel has the property that $k_{ij} = g_i h_j$ for some vectors $\mathbf{g}$ and $\mathbf{h}$. Show that convolving an image with a discrete, separable 2D filter kernel is equivalent to convolving with two 1D filter kernels. Estimate the number of operations saved for an NxN image and a $2k + 1 \times 2k + 1$ kernel.

4. Each pixel value in 500×500 pixel image $\mathcal{I}$ is an independent normally distributed random variable with zero mean and standard deviation one. Estimate the number of pixels that where the absolute value of the x derivative, estimated by forward differences (i.e. $|I_{i+1,j} - I_{i,j}|$ is greater than 3.

5. Each pixel value in 500×500 pixel image $\mathcal{I}$ is an independent normally distributed random variable with zero mean and standard deviation one. $\mathcal{I}$ is convolved with the $2k + 1 \times 2k + 1$ kernel $\mathcal{G}$. What is the covariance of pixel values in the result? (There are two ways to do this; on a case-by-case basis — e.g. at points that are greater than $2k + 1$ apart in either the x or y direction, the values are clearly independent — or in one fell swoop. Don't worry about the pixel values at the boundary.)

6. We have a camera that can produce output values that are integers in the range 0-255. Its spatial resolution is 1024 by 768 pixels, and it produces 30 frames a second. We point it at a scene that, in the absence of noise, would produce the constant value 128. The output of the camera is subject to noise that we model as zero mean stationary additive Gaussian noise with a standard deviation of 1. How long must we wait before the noise model predicts that we should see a pixel with a negative value? (Hint: you may find it helpful to use logarithms to compute the answer, as a straightforward evaluation of $\exp(-128^2/2)$ will yield 0; the trick is to get the large positive and large negative logarithms to cancel).

Practical exercises

You must do at least one of these exercises:

1. A sampled Gaussian kernel must alias, because the kernel contains components at arbitrarily high spatial frequencies. Assume that the kernel is sampled on an infinite grid. As the standard deviation gets smaller, the aliased energy must increase. Plot the energy that aliases against the standard deviation of the Gaussian kernel in pixels. Now assume that the Gaussian kernel is given on a 7x7 grid. If the aliased energy must be of the same order of magnitude as the error due to truncating the Gaussian, what is the smallest standard deviation that can be expressed on this grid?

2. Why is it necessary to check that the gradient magnitude is large at zero crossings of the Laplacian of an image? Demonstrate a series of edges for which this test is significant. The Laplacian of a Gaussian looks similar to the difference between two Gaussians at different scales. Compare these two kernels for various values of the two scales — which choices give a good approximation? How significant is the approximation error in edge finding using a zero-crossing approach?

3. Obtain an implementation of Canny's edge detector (you could try the vision home page at `http://www.somewhereorother`) and make a series of images indicating the effects of scale and contrast thresholds on the edges that are detected. How easy is it to set up the edge detector to mark only object boundaries? Can you think of applications where this would be easy? Give a set of examples that shows how effectively one can suppress texture with the scale parameter.

4. Build a normalised correlation matcher. The hand finding application is a nice one, but another may occur to you.

    - How reliable is it?
    - How many different patterns can you tell apart in practice?
    - How sensitive is it to illumination variations? shadows? occlusion?

5. Median filters can smooth corners unacceptably.

- Demonstrate this effect by applying a median filter to a variety of images; what are the qualitative effects at corners and in textured regions?
- Explain these effects.
- Show that the multi-stage median filter results in less heavily smoothed corners in practice.
- Explain this effect.
- On occasion, it is attractive to suppress texture; median filters can do this rather well. Demonstrate this phenomenon (a fluffy toy is usually a great example); why is it helpful?