Homework 2: Gaussians and PCA

CSC 2521, Fall 2004

Due: Wednesday, October 13, 3pm. This is an optional assignment — I strongly recommend that you work the problems for which the answers are not obvious.

1. Generate $N = 100$ random samples from a Gaussian distribution with mean 0 and variance 1 using MATLAB's `randn` function, and plot the histogram of these points with `hist` and an appropriate number of bins. Also, plot the Gaussian distribution; it should be approximately the same shape as the histogram (though, of course, much smoother).

   For $N = 2$ to $N = 10,000$, generate N samples from the Gaussian, and estimate the mean and variance for the value of N. Plot the estimated mean vs. N on a graph, and plot the estimated variance vs. N.

2. Derive the maximum likelihood estimate of the covariance matrix for a Gaussian distribution from measured data, in which the covariance matrix is known to be diagonal, but not necessarily spherical. In other words, the covariance matrix has the form $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1^2, \sigma_2^2, ..., \sigma_d^2)$. For convenience, you may assume that the mean of the distribution is zero.

3. Let $\mathbf{y}$ be a vector-valued variable sampled from a Probabilistic PCA model, where

$$\mathbf{x} \sim \mathcal{N}(0; \mathbf{I}) \quad (1)$$

$$\mathbf{y} \sim \mathcal{N}(\mathbf{Ax} + \mathbf{b}; \sigma^2 \mathbf{I}) \quad (2)$$

   Given a measured value of $\mathbf{y}$ and a known model ($\mathbf{A}, \mathbf{b}$, and σ^2), derive $p(\mathbf{x}|\mathbf{y}, \mathbf{A}, \mathbf{b}, \sigma^2)$, the posterior distribution over $\mathbf{x}$, using the basic rules of probability theory.

   Hint: First derive an equation for the posterior PDF. The posterior is a Gaussian distribution, but you'll have to manipulate it into the canonical form of a Gaussian with respect to $\mathbf{x}$. You can then determine what the mean and covariance of the distribution are from inspection.

4. Show that, in the PPCA model above, $p(\mathbf{y}|\mathbf{A}, \mathbf{b}, \sigma^2) = \mathcal{N}(\mathbf{y}|\mathbf{b}; \mathbf{AA}^T + \sigma^2\mathbf{I})$.

5. Using MATLAB, plot samples from a PPCA model with a 1D $\mathbf{x}$ values, and 2D $\mathbf{y}$ values. Note how the shape of the distribution changes as you vary σ^2.

6. (Optional) Are there any topics or specific papers that you are particularly interested in reading and discussing in this course? Anything you are specifically not interested in? More or less depth in advanced vision topics? Advanced learning topics without graphics applications yet?

7. (Optional) Would you prefer to only discuss 2 papers per class, thus allowing more time for detailed discussion, but covering less papers over the course of the semester?