STA 414/2104, Spring 2012, Practice Problem Set #3

Note: these problems are not for credit, and not to be handed in.

Question 1: Suppose that we are fitting a Gaussian mixture model for data items consisting of a single real value, x, using $K = 2$ components. We have $N = 5$ training cases, in which the values of x are as follows:

$$5, \quad 15, \quad 25, \quad 30, \quad 40$$

We are using the EM algorithm to find the maximum likeihood estimates for the model parameters, which are the mixing proportions for the two components, π_1 and π_2, and the means for the two components, μ_1 and μ_2. The standard deviations for the two components are fixed at 10.

Suppose that at some point in the EM algorithm, the E step found that the responsibilities of the two components for the five data items were as follows:

r_{i1}	r_{i2}
0.2	0.8
0.2	0.8
0.8	0.2
0.9	0.1
0.9	0.1

What values for the parameters π_1, π_2, μ_1, and μ_2 will be found in the next M step of the algorithm?

Question 2: Consider the factor analysis model, $x = \mu + Wz + \epsilon$, where x is an observed vector of p variables, μ is the mean vector for x, z is an unobserved vector of m common factors, W is the matrix of "factor loadings", and ϵ is a random residual. We assume that $z \sim N(0, I)$ and independently $\epsilon \sim N(0, \Sigma)$, where Σ is diagonal with diagonal entries $\sigma_1^2, \ldots, \sigma_p^2$.

Let the number of observed variables be $p = 4$ and the number of common factors be $m = 1$.

a) Give an explicit example (specifying μ, W, and Σ) showing that it is possible for the correlation of x_1 and x_2 to be negative, the correlation of x_1 and x_3 to be positive, and the correlation of x_1 and x_4 to be zero. Compute the covariance and correlation matrices of x for your example.

b) Suppose that $\mu_j = 0$ and $\sigma_j^2 = 4$ for $j = 1, 2, 3, 4$, and $W = [\,3\ 2\ 1\ 0\,]^T$. Find the covariance matrix for x, the direction of the first principal component of that covariance matrix, and the variance in that direction.

Question 3: Below is a scatterplot of 150 observations of two variables:

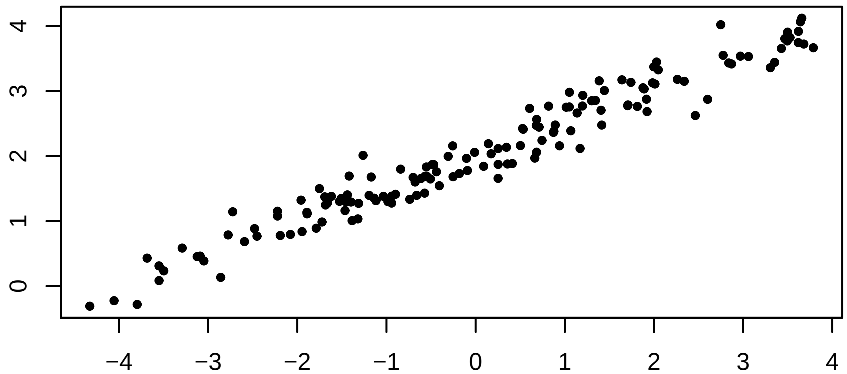

a) Write down a vector pointing in the direction of the first principal component for this data. An approximate answer found by eye is sufficient. The vector need not have length one.

   Also, draw the direction of the first principal component on the scatterplot above.

b) What is the approximate standard deviation in the first principal component's direction?

c) Suppose that each of these data points are associated with one of two classes, as shown below (with one class marked by "o" and the other by "x"):

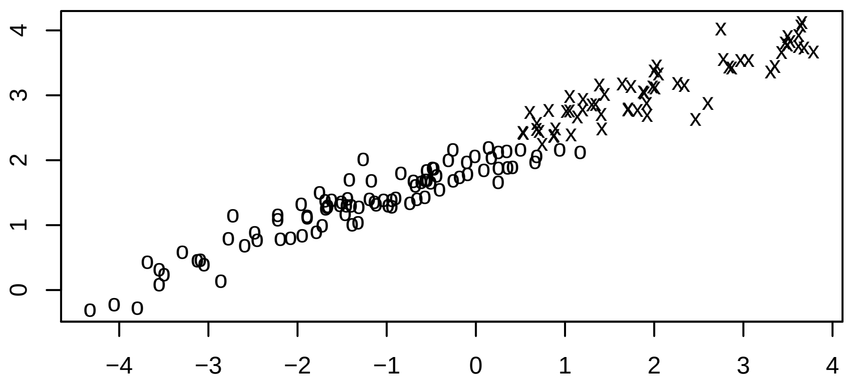

   If we reduce the data to just the projection on the first principal component, how well will we be able to classify the data points using this one number, compared to how well we would have been able to classify using the two original numbers?

d) Suppose instead that the two classes are as shown below:

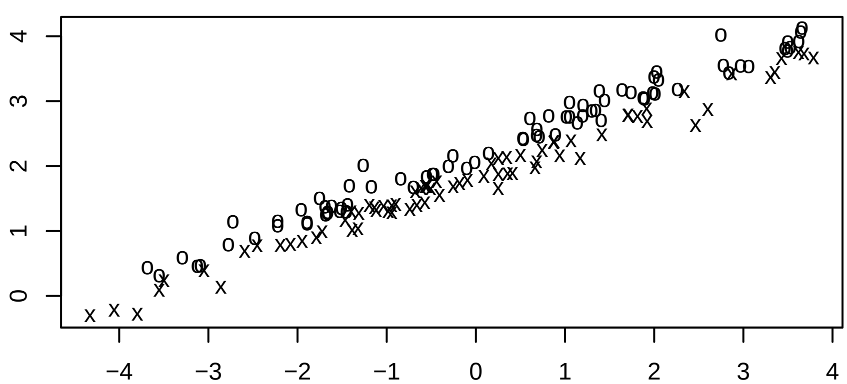

   In this case, how well will we be able to classify using just the projection on the first principal component, compared to using the two original numbers?