COMPSCI 271 - Machine Learning
Homework Due 9/27/07

1 Working with Gaussians I

Do problem 2.13 from the text. (Hint: You can use symmetry to eliminate terms and simplify your answer significantly. Your answer should contain the trace of one of the covariance matrices.)

2 Working with Gaussians II

Do problem 2.25.

3 Working with Gaussians III

Do problem 2.35. Don't forget to prove 2.124, as the rest of the problem is quite trivial. (Skip 2.291, which doesn't seem to make sense.)

4 Regression I

Do problem 3.3.

5 Regression II

Section 3.14 considers a regularized version of least squares where a penalty, λ, is applied to the square of the regression weights. a) Generalize this to the case where this is a vector of penalties $\Lambda = (\lambda_1 \dots \lambda_k)$, where λ_i is the penalty associated with w_i^2. b) Explain the relationship between the Λ and a Gaussian prior on the weights.

6 Linear Discriminants

Suppose you have constructed a linear discriminant based upon a model of the class conditional densities as Gaussians with a shared covariance matrix, Σ between both classes. Now, suppose that you are missing some features in a new datum that you must classify. Give a mathematical justification for replacing the missing features with their class conditional means. **Clarification:** Suppose your continuous features are $x_1 \dots x_n$ and that you have computed $P(x_1 \dots x_n|y)$ for each class. Now, you have a new datum where x_n is missing, so you need to evaluate $P(y|x1 \dots x_{n-1})$ for each class. Show that it's OK to pretend, when evaluating the posterior for $y = 0$, that x_n has value $E(x_n|y = 0)$, and that it's OK to pretend, when evaluating the posterior for $y = 1$, that x_n has value $E(x_n|y = 1)$. This may seem a bit surprising at first.

7 Logistic Regression

Can logistic regression classify anything that a perceptron cannot classify? Justify your answer.