UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
Department of Electrical and Computer Engineering

ECE 544NA Pattern Recognition
Fall 2007

Exam 1

Monday, October 14, 2007

- This is a CLOSED BOOK exam, but you may use ONE PAGE, BOTH SIDES of hand-written notes
- Calculators are permitted
- You must SHOW YOUR WORK to get full credit.

Problem	Score
1	
2	
3	
4	
Total	

Name: ____

Problem 1 (40 points)

The planet xkblprq and its core rotate asynchronously, in a complex multiperiodic pattern, with the result that some mornings, the sun doesn't rise. You have the following model: on the ith day, the sun rises ($x_i = 1$) with probability θ, or fails to rise ($x_i = 0$) with probability $1 - \theta$. In other words,

$$p_x(x_i) = \theta^{x_i}(1-\theta)^{(1-x_i)}$$

(a) What is $\hat{\theta}_{ML}$, the maximum likelihood estimate of θ given i.i.d. training data $\mathcal{D} = \{x_1, \ldots, x_n\}$?

(b) Prove that $s = \sum_{i=1}^{n} x_i$ is a sufficient statistic for θ.

(c) You decide to regularize your estimate with a Dirichlet prior:

$$p(\theta) = \begin{cases} \frac{(m+1)!}{k!(m-k)!}\theta^k(1-\theta)^{(m-k)} & 0 \leq \theta \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

where "!" denotes factorial. Find $\hat{\theta}_{MAP}$.

(d) Find a function which is proportional to the Bayesian estimate $\hat{p}(x_0|\mathcal{D})$, using the Dirichlet prior. In order to evaluate the integral, you may find it useful to note that the Dirichlet prior is a correctly normalized probability density.

Problem 2 (20 points)

You have landed on a planet on which it is possible, with prior probability P_1, that the sun rises every day, but it is also possible, with prior probability $1 - P_1$, that the sun only rises 50% of the time.

(a) The sun has risen for the past three days. Your commander asks you to determine, with minimum probability of error, whether or not this is a planet on which the sun always rises. For what values of P_1 would you answer "yes"?

(b) Your spaceship is expensive: each day that you sit on this planet costs \$10,000. Your goal is to determine whether or not the sun always rises, and if so, to build a \$10,000,000 solar power plant. If the sun ever fails to rise on any given day, chemicals in the solar cells will destroy the power plant, wasting \$10,000,000. As a function of P_1, how many consecutive sunrises will you watch before telling your commander that it's OK to build the plant?

Problem 3 (20 points)

You are a microcontroller with a sensor that measures the light level, x_i, in arbitrary units. Your job is to determine the weather: is today $\omega_1 =$"sunny", or $\omega_2 =$"cloudy?" You have been programmed to compute Parzen window estimates of $p(x|\omega_1)$ and $p(x|\omega_2)$ using the rectangular window:

$$\hat{p}(x) = \frac{1}{nV}\sum_{i=1}^{n}\phi\left(\frac{x - x_i}{h}\right)$$

$$\phi\left(\frac{x}{h}\right) = \begin{cases} 1 & |x| < \frac{h}{2} \\ 0 & \text{otherwise} \end{cases}$$

You have five labeled training days. Three days were sunny, with light levels of $x_1 = 4$, $x_2 = 1$, and $x_3 = 5$ units, respectively. Two days were cloudy, with light levels of $x_4 = 3$ and $x_5 = 2$ units.

(a) Plot the Parzen window estimated likelihood $\hat{p}(x|\omega_1)$ as a function of x, using $h = 1$.

(b) Plot the Parzen window estimated likelihood $\hat{p}(x|\omega_2)$ as a function of x, using $h = 1$.

(c) Plot the posterior probability, $\hat{P}(\omega_1|x)$, implied by your choices in parts (a) and (b), for all values of x for which it is defined, using $h = 1$. Assume maximum likelihood estimates of the priors $\hat{P}(\omega_1)$ and $\hat{P}(\omega_2)$.

(d) In this part, you will choose a new value of h in order to control the bias of the estimator. Suppose that you don't know anything about the true likelihood, $p(x|\omega_1)$, except that it is everywhere continuous with bounded slope:

$$\left|\frac{\partial p(x|\omega_1)}{\partial x}\right| < a$$

As in the previous sections, assume that

$$\phi\left(\frac{x}{h}\right) = \begin{cases} 1 & |x| < \frac{h}{2} \\ 0 & \text{otherwise} \end{cases}$$

Choose h, as a function of a, so that

$$|p(x|\omega_1) - E_1[\hat{p}(x|\omega_1)]| < 0.01$$

where

$$E_1[f(x)] \equiv \int f(x)p(x|\omega_1)dx$$

Problem 4 (20 points)

The dissimilarity between two d-dimensional Gaussian distributions $p(x|\omega_1)$ and $p(x|\omega_2)$, with means μ_1 and μ_2 and with covariance matrices Σ_1 and Σ_2 respectively, may be measured as

$$d(\omega_1, \omega_2) = \ln \left| \frac{1}{d^2} \text{trace}\left(\Sigma_1^{-1}\Sigma_2\right) E_1\left[(x-\mu_1)^T \Sigma_2^{-1}(x-\mu_1)\right] \right|$$

where

$$E_1\left[f(x)\right] \equiv \int f(x)p(x|\omega_1)dx$$

(a) Is $d(\omega_1, \omega_2)$ symmetric? (note: trace(AB)=trace(BA) for appropriately-dimensioned matrices.)

(b) Is $d(\omega_1, \omega_2)$ reflexive?