

Problem Set 1

MAS 622J/1.126J: Pattern Recognition and Analysis

Due Thursday, 23 September 2010

[Note: All instructions to plot data or write a program should be carried out using Matlab. In order to maintain a reasonable level of consistency and simplicity we ask that you do not use other software tools.]

If you collaborated with other members of the class, please write their names at the end of the assignment.

Problem 1: Why? [5 points]

Limit your answer to problem 1 to a page.

- a. Describe an application of pattern recognition related to your research. What are the features? What is the decision to be made? Speculate on how one might solve the problem.
- b. In the same way, describe an application of pattern recognition you would be interested in pursuing for fun in your life outside of work.

Problem 2: Probability Warm-Up [20 points]

Let X and Y be discrete random variables. Let μ_X denote the expected value of X and σ_X^2 denote the variance of X . a and b are constant values. Use excruciating detail to answer the following:

- a. Show $E[aX + bY] = aE[X] + bE[Y]$.
- b. Show $\sigma_X^2 = E[X^2] - \mu_X^2$.
- c. Show that independent implies uncorrelated.
- d. Show that uncorrelated does not imply independent.
- e. Let $Z = aX + bY$. Show that if X and Y are uncorrelated, then $\sigma_Z^2 = a^2\sigma_X^2 + b^2\sigma_Y^2$.

- f. Let X_i ($i = 1, \dots, n$) be random variables independently drawn from the same probability distribution with mean μ_X and variance σ_X^2 . For the sample mean $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$, show the following: (i) $E[\bar{X}] = \mu_X$. (ii) $\text{Var}[\bar{X}]$ (variance of the sample mean) $= \sigma_X^2/n$. Note that this is different from the sample variance $s_n^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2$.
- g. Let X_1 and X_2 be independent and identically distributed (i.i.d) continuous random variables. Can $\Pr[X_1 \leq X_2]$ be calculated? If so, find its value. If not, explain. Hint 1: Remember that for a continuous variable $\Pr[X_1 = k] = 0$, for any value of k . Hint 2: Remember the definition of i.i.d. variables.
- h. Let X_1 and X_2 be independent and identically distributed discrete random variables. Can $\Pr[X_1 \leq X_2]$ be calculated? If so, find its value. If not, explain.

Problem 3: Teatime with Gauss and Bayes [20 points]

Let $p(x, y) = \frac{1}{2\pi\alpha\beta} e^{-\left(\frac{(y-\mu)^2}{2\alpha^2} + \frac{(x-y)^2}{2\beta^2}\right)}$.

- Find $p(x)$, $p(y)$, $p(x|y)$, and $p(y|x)$. In addition, give a brief description of each of these distributions.
- Let $\mu = 0$, $\alpha = 15$, and $\beta = 3$. Plot $p(y)$ and $p(y|x = 9)$ for a reasonable range of y . What is the difference between these two distributions?

Problem 4: Covariance Matrix [15 points]

Let $\Sigma = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix}$.

- Find the eigenvalues and eigenvectors of Σ by hand (include all calculations.) Verify your computations with MATLAB function *eig*.
- Verify that Σ is a valid covariance matrix.
- We provide 200 data points sampled from the distribution $\mathcal{N}([00], \Sigma)$. Download the dataset from the course website and project the data onto the eigenvectors of the covariance matrix. What is the effect of this projection? Include the plots before and after the projection.

Problem 5: Probabilistic Modeling [20 points]

Let $x \in \{0, 1\}$ denote a person's affective state ($x = 0$ for “positive-feeling state”, and $x = 1$ for “negative-feeling state”). The person feels positive with probability θ_1 . Suppose that an affect-tagging system (or a robot) recognizes her feeling state and reports the observed state (variable y) to you. But this system is unreliable and obtains the correct result with probability θ_2 .

- Represent the joint probability distribution $P(x, y|\theta)$ for all x, y (a 2x2 matrix) as a function of the parameters $\theta = (\theta_1, \theta_2)$.
- The Maximum Likelihood estimation criterion for the parameter θ is defined as:

$$\hat{\theta}_{ML} = \arg \max_{\theta} L(t_1, \dots, t_n; \theta) = \arg \max_{\theta} \prod_{i=1}^n p(t_i|\theta)$$

where we have assumed that each data point t_i is drawn independently from the same distribution so that the likelihood of the data is $L(t_1, \dots, t_n; \theta) = \prod_{i=1}^n p(t_i|\theta)$. Likelihood is viewed as a function of the parameters, which depends on the data. Since the above expression can be technically challenging, we maximize the log-likelihood $\log L(t_1, \dots, t_n; \theta)$ instead of likelihood. Note that any monotonically increasing function (i.e., log function) of the likelihood has the same maxima. Thus,

$$\hat{\theta}_{ML} = \arg \max_{\theta} \log L(t_1, \dots, t_n; \theta) = \arg \max_{\theta} \sum_{i=1}^n \log p(t_i|\theta)$$

Suppose we get the following joint observations $t = (x, y)$.

x	y
1	0
1	1
0	0
1	1
0	0
0	1
0	0
1	1

What are the maximum-likelihood (ML) values of θ_1 and θ_2 ? (*Hint.* Since $P(x, y|\theta) = P(y|x, \theta_2)P(x|\theta_1)$, the estimation of the two parameters can be done separately in the log-likelihood criterion.)

Problem 6: Ring Problem [20 points]

To get credit for this problem, you must not only write your own correct solution, but also write a computer simulation (in either Matlab or Python) of the process of playing this game:

Suppose I hide the ring of power in one of three identical boxes while you weren't looking. The other two boxes remain empty. After hiding the ring of power, I ask you to guess which box it's in. I know which box it's in and, after you've made your guess, I deliberately open the lid of an empty box, which is one of the two boxes you did not choose. Thus, the ring of power is either in the box you chose or the remaining closed box you did not choose. Once you have made your initial choice and I've revealed to you an empty box, I then give you the opportunity to change your mind – you can either stick with your original choice, or choose the unopened box. You get to keep the contents of whichever box you finally decide upon.

- What choice should you make in order to maximize your chances of receiving the ring of power? Justify your answer using Bayes' rule.
- Write a simulation. There are two choices in this game for the contestant in this game: (1) choice of box, (2) choice of whether or not to switch. In your simulation, first let the host choose a random box to place the ring of power. Show a trace of your program's output for a single game play, as well as a cumulative probability of winning for 1000 rounds of the two policies (1) to choose a random box and then switch and (2) to choose a random box and not switch.