

CSE 555 Spring 2009 Homework 1: Bayesian Decision Theory

Jason J. Corso

Computer Science and Engineering

University at Buffalo SUNY

jcorso@cse.buffalo.edu

Date Assigned 18 Jan 2009

Date Due 3 Feb 2009

Homework must be submitted in class. No late work will be accepted.

Problem 1: Bayesian Decision Rule (25%)

Suppose the task is to classify the input signal x into one of K classes $\omega \in \{1, 2, \dots, K\}$ such that the action $\alpha(x) = i$ means classifying x into class i . The Bayesian decision rule is to maximize the posterior probability

$$\alpha_{\text{Bayes}}(x) = \omega^* = \arg \max_{\omega} p(\omega|x) .$$

Suppose we replace it by a *randomized decision rule*, which classifies x to class i following the posterior probability $p(\omega = i|x)$, i.e.,

$$\alpha_{\text{rand}}(x) = \omega \sim p(\omega|x) .$$

1. What is the overall risk R_{rand} for this decision rule? Derive it in terms of the posterior probability using the zero-one loss function.
2. Show that this risk R_{rand} is always no smaller than the Bayes risk R_{Bayes} . Thus, we cannot benefit from the randomized decision.
3. Under what conditions on the posterior are the two decision rules the same?

Problem 2: Bayesian Classification Boundary (25%)

For a two-class recognition problem with salmon ($\omega = 1$) and sea bass ($\omega = 2$), suppose we have two features $\mathbf{x} = (x_1, x_2)$ and the two class-conditional densities, $p(x|\omega = 1)$ and $p(x|\omega = 2)$, are 2D Gaussian distributions centered at points $(4, 16)$ and $(16, 4)$ respectively with the same covariance matrix $\Sigma = 4\mathbf{I}$ (with \mathbf{I} is the identity matrix). Suppose the priors are $P(\omega = 1) = 0.6$ and $P(\omega = 2) = 0.4$.

1. Suppose we use a Bayes decision rule, write the two discriminant functions $g_1(\mathbf{x})$ and $g_2(\mathbf{x})$.
2. Derive the equation for the decision boundary $g_1(\mathbf{x}) = g_2(\mathbf{x})$. Draw the boundary on the feature space (the 2D plane).

Problem 3: Working with the Normal Distribution (10%)

This is part of Problem #18 in DHS Chapter 2.

Consider two normal distributions in one dimension: $N(\mu_1, \sigma_1^2)$ and $N(\mu_2, \sigma_2^2)$. Imagine that we choose two random samples x_1 and x_2 , one from each of the two distributions and calculate their sum $x_3 = x_1 + x_2$. Suppose we do this repeatedly.

1. Consider the resulting distribution of values of x_3 . Show that x_3 possesses the requisite statistical properties and thus its distribution is normal.
2. What is the mean, μ_3 , of the new distribution?
3. What is the variance, σ_3^2 , of the new distribution?

Problem 4: Bayesian Reasoning (40%)

Formulate and solve this classical problem using the Bayes rule. There are three criminals A, B, and C waiting in three separate jail cells. One of them will be executed in the next morning when the sun rises. A is very nervous, as he has 1/3 chance to be the one. He tries to get some information from the janitor: "I know you cannot tell me whether I will be executed in the next morning, but can you tell me which of my inmates B and C will not be executed? Because one of them will not be executed anyway, by pointing out who will not be executed, you are not telling me any information." This sounds quite logical. So the Janitor tells A that C won't be executed. At a second thought, A gets much more worried. Before he asked the janitor, he thought he had 1/3 chance, but with C excluded, he seems to have 1/2 chance. A says to himself: "What did I do wrong? Why did I ask the janitor?"

1. Formulate the problem using the Bayes rule, i.e. what are the random variables and the input data? What are the meaning of the prior and posterior probabilities in this problem?
2. What are the probability values for the prior?
3. What are the probability values for the likelihood?
4. Calculate the posterior probability (you need to derive the probability values with intermediate steps, not simply showing the final values).
5. What is the probability of A being executed after he knows that C is excluded?
6. Did the janitor tell us any information about A's fate?
7. Explain how the Bayes rule helps you.