Decision theory

Introduction to Machine Learning (CSCI 1950-F), Summer 2011

Solve the following problems. Provide mathematical justification for your answers.

(1) The reject option

Consider a binary classification problem where you are given $p(y|x) = P(Y = y|X = x)$, the probability of class $y \in \{0, 1\}$ given the point x. Suppose that the classifier function $f(x)$ is allowed to take one of 3 values: 0, 1, or *reject*. These correspond to the decisions: choose class 0, choose class 1, or "reject" x (that is, refuse to make a decision). (This type of setup is useful when we want a classifier to be able to say, in effect, "I don't know".)

Assume the following loss function:

$$L(y, \hat{y}) = \begin{cases} 0 & \text{if } \hat{y} = y \\ 1 & \text{if } \hat{y} \neq y \text{ and } \hat{y} \in \{0, 1\} \\ \lambda & \text{if } \hat{y} = \mathit{reject} \end{cases}$$

where $\lambda > 0$. Find the optimal decision rule $f(x)$ to minimize the (conditional) expected loss $E(L(Y, f(x))|X = x)$. Express your answer in terms of $p(y|x)$ and λ. Simplify your answer as much as possible.

(2) Absolute (or L^1) loss

For a real-valued random variable Y with a density, perhaps the most common loss function is the square (or L^2) loss, $L(y, \hat{y}) = (y - \hat{y})^2$, however in some applications it will be inappropriate since penalizing the error quadratically makes the resulting decision rule very sensitive to outliers. (For example, suppose you are getting GPS measurements of your location, but your GPS device occasionally malfunctions and reports a randomly chosen location on Earth. A problem like this actually occurs in robotics.) A loss function which is more robust to outliers is the absolute (or L^1) loss: $L(y, \hat{y}) = |y - \hat{y}|$, that is, just the distance between y and $\hat{y}$.

Assume Y has a density $p(y)$, and assume that this density function is continuous. Show that any median of the distribution of Y minimizes the expected loss $EL(Y, a)$ (as a function of a) when L is the absolute loss. (Here, since Y has a density, a median is a real number m such that $P(Y \leq m) = P(Y \geq m) = 1/2$.)

Try to solve this on your own before looking at the hint below.

(Hint: Do the usual calculus thing. The 1st Fundamental Theorem of Calculus may be helpful. Also, if it helps, you can assume that it's okay to "differentiate under the integral sign".)