STAT3007/7007 Deep Learning, Tutorial 3

2022 Semester 2

1. (Bayes optimal classifier)

   (a) If the probabilities that tomorrow will be rainy, sunny, or cloudy are 0.2, 0.3 and 0.5 respectively, what is the weather condition that you will predict to minimize your expected prediction error?

   (b) Consider the problem of classifying a real-valued input X to a class $Y \in \{0, 1\}$. Suppose the true data distribution $P(X, Y) = P(Y)P(X \mid Y)$ is defined by

   $$P(Y = 0) = P(Y = 1) = \frac{1}{2}, \tag{1}$$
   $$P(X \mid Y = 0) = N(X; 0, 1), \tag{2}$$
   $$P(X \mid Y = 1) = N(X; 1, 1), \tag{3}$$

   where $N(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$ is the probability density function (PDF) for the normal distribution $N(\mu, \sigma^2)$.

   Draw the decision boundary for the Bayes classifier. Justify your answer.

2. (Regression function) We are given a feature vector $\mathbf{x}$, and we want to predict a real-valued output y. Assume that the conditional distribution $P(Y \mid \mathbf{x})$ is given by

y	0.1	0.2	0.3	0.4	0.5
$P(y \mid \mathbf{x})$	0.1	0.1	0.1	0.2	0.5

   (a) What output value $\hat{y} \in \mathbf{R}$ should we predict if we want to minimize the expected quadratic loss $\mathbb{E}_{Y \mid \mathbf{x}}(Y - \hat{y})^2$?

   (b) What output value $\hat{y} \in \mathbf{R}$ should we predict if we want to minimize the expected ℓ_1 loss $\mathbb{E}_{Y \mid \mathbf{x}} |Y - \hat{y}|$?