COMPUTER SCIENCE TRIPOS Part II – 2020 – Paper 8

10 Machine Learning and Bayesian Inference (sbh11)

(*a*) Give a detailed description of the general *Bayes decision rule* for classification. Include in your answer definitions of the *loss*, *conditional risk*, *decision rule* and *risk*. [7 marks]

(*b*) For a problem with C classes, we suffer a loss of 1 for an incorrect classification and 0 for a correct one. Show that the Bayes decision rule for inputs $\mathbf{x}$ is

$$h(\mathbf{x}) = \text{argmax}_c \Pr(c|\mathbf{x}).$$

[3 marks]

(*c*) For a problem with 2 classes, we now have three possibilities: classify $\mathbf{x}$ as being in class c_1, classify $\mathbf{x}$ as being in class c_2, or decline to classify $\mathbf{x}$. Classifying some $\mathbf{x}$ correctly results in a loss of 0 and classifying it incorrectly results in a loss of 1. Declining to classify $\mathbf{x}$ has a cost of θ_1 if $\mathbf{x}$ should be classified in class c_1 and θ_2 if it should be classified in class c_2. Both θ_1 and θ_2 can take values between 0 and 1/2.

Give a graphical representation of the conditional risks and use it to show that the Bayes decision rule for this problem is:

$$h(\mathbf{x}) = \begin{cases} c_1 & \text{if } p \leq q_1 \\ \text{Decline} & \text{if } q_1 < p \leq q_2 \\ c_2 & \text{if } q_2 < p \end{cases}$$

where $p = \Pr(c_1|\mathbf{x})$, $q_1 = \frac{\theta_2}{1-(\theta_1-\theta_2)}$ and $q_2 = \frac{1-\theta_2}{1+(\theta_1-\theta_2)}$. [10 marks]