COMPUTER SCIENCE TRIPOS Part II – 2025 – Paper 9

8 Machine Learning and Bayesian Inference (sbh11)

You have a two-class classification problem, with classes $\{c_+, c_-\}$. Rather than making a simple prediction of which class an example $\mathbf{x}$ falls into, you wish to make two kinds of judgement:

(a) $\mathbf{x}$ is predicted to be in class c_i, where $i \in \{+, -\}$.

(b) $\mathbf{x}$ is *weakly* predicted to be in class c_i, where $i \in \{+, -\}$.

Predicting class membership according to (a) results in a loss of 0 if the prediction is correct and a loss of 1 if in error. Weakly predicting class membership according to (b) results in a loss of θ_1 if correct and a loss of $1 - \theta_2$ if in error. We assume $0 \leq \theta_i \leq 1$ where $i \in \{1, 2\}$. Your aim is to design a *Bayes decision rule* for this problem.

(a) Denoting by $\Pr(C|\mathbf{x})$ the conditional distribution of the class, give a definition of *conditional risk*. [3 marks]

(b) Let $p = \Pr(c_+|\mathbf{x})$. Write down expressions for the *conditional risks* for the actions described. [4 marks]

(c) What constraint should be placed on the values θ_1 and θ_2 such that the options for weak predictions will be relevant in applying the Bayes decision rule? [4 marks]

(d) Derive the Bayes decision rule for this problem, assuming that the constraints derived in Part (c) are met. [4 marks]

You now wish to add a third possibility: 'I *decline* to make a prediction for $\mathbf{x}$'. Declining to make a prediction always has a loss of θ_3, where $0 \leq \theta_3 \leq 1$.

(e) What additional constraint is needed on θ_3 such that it will be relevant in applying the Bayes decision rule? [2 marks]

(f) What modification is necessary to your answer to Part (d) in order to include the option to decline in the Bayes decision rule? [3 marks]