Lecture 4: Quiz (Bayesian Tests and Model Selection)

Instructor: Thibault Randrianarisoa

Questions

Question 1 Regarding the Bayes test for a 0-1 (balanced) loss function, which of the following statements is true?

a) The test minimizes the Type I error regardless of the Type II error.

b) The test rejects the null hypothesis H_0 if the posterior probability $\pi(\Theta_0|X)$ is less than or equal to $\frac{1}{2}$.

c) The test compares the likelihood ratio to a fixed constant α determined by the user.

d) The test accepts the alternative hypothesis H_1 if the posterior probability $\pi(\Theta_1|X)$ is greater than $\pi(\Theta_0|X)$.

Question 2 In a weighted loss function scenario where a_0 is the cost of a Type I error and a_1 is the cost of a Type II error, what is the decision if $a_0 = 0$?

a) Always reject H_1

b) Always reject H_0

c) Reject H_0 if $\pi(\Theta_0|X) \leq 1/2$

d) Accept H_0 if $\pi(\Theta_0|X) \geq 1/2$

Question 3 Consider testing a point null hypothesis $H_0 : \theta = \theta_0$ against a composite alternative $H_1 : \theta \neq \theta_0$. If one incorrectly employs a continuous prior density $\pi(\theta)$ (such as a Gaussian) over the entire parameter space $\mathbb{R}$ without assigning specific mass to θ_0, what is the unavoidable consequence for the Bayesian test?

a) The Bayes Factor will always equal 1, rendering the test inconclusive regardless of the data.

b) The posterior probability $\pi(\Theta_0|X)$ will be 0, leading to the guaranteed rejection of H_0 even if the null is true.

c) The test will behave identical to a Frequentist Likelihood Ratio test.

d) The marginal likelihood will become infinite, causing the posterior distribution to be improper.

Question 4 For simple hypotheses H_0 and H_1 with equal priors, the minimum Bayes Risk is related to the Total Variation distance (d_{TV}) by $R_B(\pi) = \frac{1}{2}(1 - d_{TV}(P_0, P_1))$. Based on this relationship, which scenario represents the **highest** possible risk (worst performance) for the decision maker?

a) When the distributions P_0 and P_1 are completely disjoint (support intersection is empty).

b) When the affinity $\mathcal{A}(P_0, P_1)$ between the distributions is 0.

c) When the distributions are identical ($P_0 = P_1$).

d) When the Total Variation distance $d_{TV}(P_0, P_1)$ approaches 1.

Solutions

1. Correct Answer: b) The test rejects the null hypothesis H_0 if...
Reasoning: The Bayes test for a 0-1 loss function is given by $\varphi^*(X) = \mathbb{I}_{\pi(\Theta_\circ|X)\leq\frac{1}{2}}$. This means we reject H_0 (set $\varphi = 1$) when the posterior probability of Θ_0 is less than or equal to the posterior probability of Θ_1 (which sums to 1, hence the 1/2 threshold).

2. Correct Answer: b) Always reject H_0
Reasoning: We check if $\pi(\Theta_0|X) \leq 1$, which is always true. Since Type I error has no cost, we can minimize the Type II error by always rejecting the null.

3. Correct Answer: b) The posterior probability $\pi(\Theta_0|X)$ will be 0...
Reasoning: A continuous prior assigns zero probability mass to any single point, including θ_0. Since $\pi(\Theta_0) = \int_{\{\theta_0\}} \pi(\theta)d\theta = 0$, the posterior probability $\pi(\Theta_0|X)$ will also be zero (as long as the likelihood is finite). This forces the rejection of H_0 regardless of the evidence.

4. Correct Answer: c) When the distributions are identical ($P_0 = P_1$).
Reasoning: The risk is maximized when the distributions are hardest to distinguish.

- If $P_0 = P_1$, then $d_{TV}(P_0, P_1) = 0$.
- Plugging this into the formula: $R_B(\pi) = \frac{1}{2}(1 - 0) = 0.5$.

This represents a 50% chance of error (random guessing), which is the worst possible performance in a binary decision. Conversely, if $d_{TV} = 1$ (disjoint), the risk drops to 0.