STAD91: Quiz 1 Practice

Instructor: Thibault Randrianarisoa

Questions

Question 1 You have designed a new estimator $\hat{\theta}_n$. You have successfully proven that its quadratic risk converges to zero as $n \to \infty$ (i.e., $\mathbb{E}[(\hat{\theta}_n - \theta)^2] \to 0$). Which form of convergence is **NOT** automatically guaranteed?

a) Convergence in Probability ($\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta$).

b) Convergence in Distribution ($\hat{\theta}_n \xrightarrow{\mathcal{L}} \theta$).

c) Almost Sure Convergence ($\hat{\theta}_n \xrightarrow{a.s.} \theta$).

d) Convergence in L^2.

Question 2 You are working with a Bernoulli(θ) model where the Maximum Likelihood Estimator is the empirical mean, $\hat{\theta}_n = \overline{X}_n$. This estimator is asymptotically normal with asymptotic variance $\theta(1-\theta)$. What is a $1-\alpha$ asymptotic confidence interval for θ?

a) $\left[\overline{X}_n \pm \frac{q_\alpha}{\sqrt{n}}\right]$

b) $\left[\overline{X}_n \pm \frac{q_\alpha\sqrt{\theta(1-\theta)}}{\sqrt{n}}\right]$

c) $\left[\overline{X}_n \pm \frac{q_\alpha\sqrt{\overline{X}_n(1-\overline{X}_n)}}{\sqrt{n}}\right]$

d) $\left[\overline{X}_n \pm \frac{q_\alpha\overline{X}_n(1-\overline{X}_n)}{\sqrt{n}}\right]$

Question 3 Based on the properties of distributions defined in the slides, which of the following random variables does **NOT** follow a Uniform distribution on $[0, 1]$?

a) A random variable $X \sim Beta(1, 1)$.

b) The ratio $R = \frac{E_1+E_2}{E_1+E_2+E_3+E_4}$, where E_i are independent $\Gamma(1/2, \lambda)$ variables.

c) The first component X_1 of a random vector $(X_1, X_2) \sim Dir(1, 1)$.

d) The first component X_1 of a random vector $(X_1, X_2, X_3) \sim Dir(1, 1, 1)$.

Question 4 If you choose to use a different, suboptimal predictor $\varphi(X)$ instead of the optimal $m(X)$ to minimize $\mathbb{E}[(Y - \varphi(X))^2]$, what exactly is the extra penalty you incur?

a) The penalty is the cross-term: $2\mathbb{E}[(Y - m(X))(m(X) - \varphi(X))]$.

b) The penalty is the Mean Squared Error of the optimal predictor itself: $\mathbb{E}[(Y - m(X))^2]$.

c) The penalty is the mean squared distance between your predictor and the optimal one: $\mathbb{E}[(m(X) - \varphi(X))^2]$.

d) The penalty is zero as long as your predictor $\varphi(X)$ is unbiased.

Solutions

1. Correct Answer: c) Almost Sure Convergence.
Reasoning: Convergence of the Quadratic Risk to zero is equivalent to Convergence in L^2. L^2 convergence implies Convergence in Probability, which in turn implies Convergence in Distribution. However, there is no implication from L^2 or Probability to Almost Sure convergence.

2. Correct Answer: c) $\left[\overline{X}_n \pm \frac{q_\alpha \sqrt{\overline{X}_n(1-\overline{X}_n)}}{\sqrt{n}}\right]$
Reasoning: The general construction requires replacing the unknown true variance $\sigma^2(\theta)$ with the estimated variance $\sigma^2(\hat{\theta}_n)$. Since the asymptotic variance is $\theta(1-\theta)$, the estimated standard deviation is $\sqrt{\overline{X}_n(1-\overline{X}_n)}$. Answer (b) is incorrect because it relies on the unknown θ.

3. Correct Answer: d) The first component X_1 of $(X_1, X_2, X_3) \sim Dir(1,1,1)$.
Reasoning:

- a) Slides state explicitly $Beta(1,1) = \mathcal{U}[0,1]$.
- b) By Gamma additivity, $E_1 + E_2 \sim \Gamma(1,\lambda)$ and $E_3 + E_4 \sim \Gamma(1,\lambda)$, independently. The ratio of independent $\Gamma(a,\lambda)$ and $\Gamma(b,\lambda)$ is $Beta(a,b)$. Here $a = 1, b = 1$, so it is $Beta(1,1)$ (Uniform).
- c) For $K = 2$, $Dir(\alpha_1, \alpha_2)$ is $Beta(\alpha_1, \alpha_2)$. So this is $Beta(1,1)$ (Uniform).
- d) Marginals are $X_i \sim Beta(\alpha_i, \sum \alpha_k - \alpha_i)$. Here $\alpha_1 = 1$ and $\sum \alpha_k = 3$. So $X_1 \sim Beta(1,2)$, which is **not** Uniform.

4. Correct Answer: c) The penalty is the mean squared distance between your predictor and the optimal one: $\mathbb{E}[(m(X) - \varphi(X))^2]$.
Reasoning: The error decomposes into the irreducible error plus a non-negative term: $\mathbb{E}[(\mathbb{E}[Y|X] - \varphi(X))^2]$. The cross-term (Answer a) vanishes (is zero) due to the orthogonality property.