CSI973 Final
Take-Home
Due May 11, 2004

This is not meant to be an endurance test!

This exam is to be completed without communication with any person (other than perhaps the instructor). You may use any book or any notes from class. Be sure to write clearly and to show all steps.

Notation for distributions usually follows that of Shao, but sometimes it's closer to what I used in the handouts.

1. Let X and Y be independent random variables with $X \sim \mathrm{N}(\mu_X, 1)$ and $Y \sim \mathrm{N}(\mu_Y, 1)$. Consider testing the simple hypothesis $H_0 : \mu_X = \mu_Y = 0$ versus the one-sided alternative $H_a : \mu_X > 0,\ \mu_Y > 0$. Show that a UMPU test does not exist (although this is an exponential family).

2. (a) Consider the Cauchy family, with density function

$$p_C(x;\theta) = \frac{1}{\pi(1 + (x-\theta)^2)}.$$

   Show that the likelihood ratio is not monotone, and show that no UMP one-sided test exists.

   (b) Now consider the uniform family, with density function

$$p_U(x;\theta) = 1 \text{ for } x \in (\theta, \theta + 1),$$

   and zero elsewhere.

   Show that the likelihood ratio is not monotone. Derive a UMP one-sided test for $H_0 : \theta = \theta_0$ versus $H_a : \theta > \theta_0$.

3. Consider a sample $X_1, \ldots, X_n$ from the discrete uniform distribution with probability mass $1/N$ on each of the integers $1, \ldots, N$. Consider testing $H_0 : N = N_0$ versus $H_a : N = N_1$, where $N_1 > N_0$.

   (a) Develop a randomized UMP test based on the likelihood ratio.

   (b) Show that there are only two possible sizes for a nonrandomized UMP test based on the likelihood ratio. (What are the sizes?)

4. Consider the problem of comparing regression lines for m sets of observations (Y_{ij}, z_{ij}), the i^{th} set of which contains n_i observations. Assume the model for $\mathrm{E}(Y_{ij} = \mu_{ij})$

$$\mu_{ij} = \alpha_i + \beta_i z_{ij}, \quad j = 1, \ldots, n_i, \quad i = 1, \ldots, m.$$

   Develop a test of the hypothesis $H_0 : \beta_1 = \cdots = \beta_m$ against the alternative that at least two are different from each other. Discuss the properties of your test.

5. Let $X_1, \ldots, X_n$ be independently distributed from a location family with means $\mu_1, \ldots, \mu_n$ and common finite variance σ^2. We assume that for each n, the vector $(\mu_1, \ldots, \mu_n)$ lies in an s-dimensional subspace of $\mathbb{R}^n$ with s fixed. Test statistics for r restrictions on the means in this context have the form

$$\frac{\left(\sum_{i=1}^n (X_i - \hat{\hat{\mu}}_i)^2 - \sum_{i=1}^n (X_i - \hat{\mu}_i)^2\right)/r}{\sum_{i=1}^n (X_i - \hat{\mu}_i)^2/(n-s)},$$

where $\hat{\hat{\mu}}_i$ is the least squares estimator of μ_i under the restrictions, and $\hat{\mu}_i$ is the least squares estimator of μ_i with no restrictions.

Show that the denominator of this expression converges in probability to σ^2 as $n \to \infty$.

6. Let $\Theta = (M, \Sigma^2)$ (here M and Σ^2 are scalars). Let $X_1, \ldots, X_n$ given $\Theta = (\mu, \sigma^2)$ be iid $\mathrm{N}(\mu, \sigma^2)$. We want to test $H_0 : M = \mu_0$. For $M \neq \mu_0$, we assume a prior distribution on $1/\Sigma^2$ of gamma with parameters $a_0/2$ and $b_0/2$. (This is an "inverse gamma" distribution. The "2" is used for simplicity.) For given $\Sigma^2 = \sigma^2$, we assume the prior distribution on M to be $\mathrm{N}(\xi, \sigma^2/\lambda_0)$. (This is the usual conjugate prior.) Conditional on $M = \mu_0$, using the above, we have a prior distribution on Σ^2 to be an inverse gamma with parameters $a_0^*/2$ and $b_0^*/2$, where $a_0^* = a_0 + 1$ and $b_0^* = b_0 + \lambda_0(\mu_0 - \xi)^2$.

   (a) Work out the conditional probability density of (X, Σ) given $M = \mu_0$.

   (b) Work out the conditional probability density of (X, Σ, M) given $M \neq \mu_0$.

   (c) Now work out the Bayes factor for the test (just integrate out the parameters and take the ratio).