MTL108: Solution to Problem Set-12

IIT Delhi

Notation: Throughout these solutions, we will consistently use uppercase letters (e.g., X_i, Y_i) to denote the random variables in our sample.

Problem 1: Testing the Mean of a Normal Population

Let $X_1, X_2, \ldots, X_n \sim N(\mu, \sigma^2)$ be an independent and identically distributed (i.i.d.) random sample.

(a) **Known Variance ($\sigma^2 = \sigma_0^2$):**

We are testing the simple null hypothesis $H_0 : \mu = \mu_0$ against the simple alternative $H_1 : \mu = \mu_1$ (assuming $\mu_1 > \mu_0$).

To find the most powerful test, we use the Neyman-Pearson Lemma. The joint likelihood function for the normal sample is:

$$L(\mu) = (2\pi\sigma_0^2)^{-n/2} \exp\left(-\frac{1}{2\sigma_0^2}\sum_{i=1}^{n}(X_i - \mu)^2\right)$$

The lemma dictates that we should reject H_0 in favor of H_1 if the likelihood ratio $\lambda(X) = \frac{L(\mu_1)}{L(\mu_0)}$ is sufficiently large (greater than some constant k):

$$\frac{\exp\left(-\frac{1}{2\sigma_0^2}\sum(X_i - \mu_1)^2\right)}{\exp\left(-\frac{1}{2\sigma_0^2}\sum(X_i - \mu_0)^2\right)} \geq k$$

Let's simplify this by taking the natural logarithm of both sides. Because the logarithm is a strictly increasing function, it preserves the direction of the inequality:

$$-\frac{1}{2\sigma_0^2}\left[\sum_{i=1}^{n}(X_i - \mu_1)^2 - \sum_{i=1}^{n}(X_i - \mu_0)^2\right] \geq \ln k$$

Expanding the squared terms inside the bracket:

$$-\frac{1}{2\sigma_0^2}\left[-2\mu_1\sum_{i=1}^{n}X_i + n\mu_1^2 + 2\mu_0\sum_{i=1}^{n}X_i - n\mu_0^2\right] \geq \ln k$$

Grouping the $\sum X_i$ terms together yields:

$$\frac{\mu_1 - \mu_0}{\sigma_0^2}\sum_{i=1}^{n}X_i \geq \ln k + \frac{n(\mu_1^2 - \mu_0^2)}{2\sigma_0^2}$$

Because we specified that $\mu_1 > \mu_0$, the coefficient $\frac{\mu_1 - \mu_0}{\sigma_0^2}$ is strictly positive. Therefore, when we divide both sides by this coefficient, the inequality sign does not flip. We can absorb all the constants on the right side into a new critical threshold, c:

$$\bar{X} \geq c$$

Probabilistic Distribution & Rejection Region: Intuitively, we reject H_0 if the sample mean is unusually large. To find the exact critical value, we look at the distribution of $\bar{X}$ under the null hypothesis. Since linear combinations of normal variables are normal, $\bar{X} \sim N(\mu_0, \sigma_0^2/n)$.

Standardizing this gives us a test statistic that perfectly follows a Standard Normal distribution:

$$Z = \frac{\bar{X} - \mu_0}{\sigma_0/\sqrt{n}} \sim N(0, 1)$$

Thus, our rejection rule at significance level α becomes: Reject H_0 if $Z \geq z_\alpha$.

Conclusion: Notice that our final test statistic and critical region rely only on α, μ_0, and σ_0. The specific value of μ_1 dropped out entirely! Because this single optimal test works for *any* $\mu_1 > \mu_0$, it forms a Most Powerful (MP) test for the composite alternative $H_1 : \mu > \mu_0$.

(b) **Unknown Variance:**

Now we test $H_0 : \mu = \mu_0$ against the two-sided alternative $H_1 : \mu \neq \mu_0$. Because the variance is unknown, our parameter space is two-dimensional: $\theta = (\mu, \sigma^2)$. The Neyman-Pearson lemma no longer applies, so we use the Likelihood Ratio Test (LRT).

The unrestricted Maximum Likelihood Estimates (MLEs) for the whole parameter space are $\hat{\mu} = \bar{X}$ and $\hat{\sigma}^2 = \frac{1}{n}\sum(X_i - \bar{X})^2$. Under the restriction of H_0, the MLEs become $\hat{\mu}_0 = \mu_0$ and $\hat{\sigma}_0^2 = \frac{1}{n}\sum(X_i - \mu_0)^2$.

The Likelihood Ratio (Λ) evaluates the ratio of the restricted likelihood to the unrestricted likelihood:

$$\Lambda = \frac{L(\hat{\mu}_0, \hat{\sigma}_0^2)}{L(\hat{\mu}, \hat{\sigma}^2)} = \frac{(2\pi\hat{\sigma}_0^2)^{-n/2} e^{-n/2}}{(2\pi\hat{\sigma}^2)^{-n/2} e^{-n/2}} = \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}\right)^{n/2} = \left(\frac{\sum(X_i - \bar{X})^2}{\sum(X_i - \mu_0)^2}\right)^{n/2}$$

We can decompose the denominator using the algebraic identity $\sum(X_i - \mu_0)^2 = \sum(X_i - \bar{X})^2 + n(\bar{X} - \mu_0)^2$ (where the cross-term sums to zero). Substituting this back into Λ:

$$\Lambda = \left(\frac{1}{1 + \frac{n(\bar{X} - \mu_0)^2}{\sum(X_i - \bar{X})^2}}\right)^{n/2} = \left(1 + \frac{T^2}{n-1}\right)^{-n/2}$$

where $T = \frac{\bar{X} - \mu_0}{S/\sqrt{n}}$ and the sample variance is $S^2 = \frac{1}{n-1}\sum(X_i - \bar{X})^2$.

The LRT states we should reject H_0 for small values of Λ. Looking at the equation above, Λ shrinks as T^2 grows. Therefore, rejecting for small Λ is mathematically identical to rejecting for large values of $|T|$.

Probabilistic Distribution & Rejection Region: Because we are estimating the unknown population variance using S^2, our test statistic T incorporates additional uncertainty. Instead of a Standard Normal, T follows a Student's t-distribution with $n - 1$ degrees of freedom (accounting for the heavier tails). Under H_0, we reject if $|T| \geq t_{\alpha/2, n-1}$.

Problem 2: Testing the Variance of a Normal Population

Let $X_1, X_2, \ldots, X_n \sim N(\mu, \sigma^2)$.

(a) **Known Mean ($\mu = \mu_0$):**

We test $H_0 : \sigma^2 = \sigma_0^2$ against $H_1 : \sigma^2 = \sigma_1^2$ (assuming $\sigma_1^2 > \sigma_0^2$). Using the Neyman-Pearson Lemma, we set up the likelihood ratio:

$$\frac{L(\sigma_1^2)}{L(\sigma_0^2)} = \frac{(\sigma_1^2)^{-n/2} \exp\left(-\frac{1}{2\sigma_1^2} \sum (X_i - \mu_0)^2\right)}{(\sigma_0^2)^{-n/2} \exp\left(-\frac{1}{2\sigma_0^2} \sum (X_i - \mu_0)^2\right)} \geq k$$

Taking the natural log to isolate the sum:

$$-\frac{n}{2} \ln\left(\frac{\sigma_1^2}{\sigma_0^2}\right) + \frac{1}{2}\left(\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right) \sum_{i=1}^{n} (X_i - \mu_0)^2 \geq \ln k$$

Since we assumed $\sigma_1^2 > \sigma_0^2$, the factor $\left(\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right)$ is strictly positive. We can safely rearrange the inequality without flipping the sign, isolating the sum of squared deviations:

$$\sum_{i=1}^{n} (X_i - \mu_0)^2 \geq c$$

Probabilistic Distribution & Rejection Region: Recall that standardizing a normal variable and squaring it yields a Chi-squared distribution with 1 degree of freedom: $\left(\frac{X_i - \mu_0}{\sigma_0}\right)^2 \sim \chi_1^2$. Because our sample is independent, the sum of n such variables gives our test statistic Q:

$$Q = \sum_{i=1}^{n} \left(\frac{X_i - \mu_0}{\sigma_0}\right)^2 \sim \chi_n^2$$

We reject H_0 if $Q \geq \chi^2_{\alpha, n}$.

(b) **Unknown Mean:**

We now test $H_0 : \sigma^2 = \sigma_0^2$ against $H_1 : \sigma^2 \neq \sigma_0^2$ using the LRT. The unrestricted MLEs are $\hat{\mu} = \bar{X}$ and $\hat{\sigma}^2 = \frac{1}{n} \sum (X_i - \bar{X})^2$. The restricted MLEs (forcing the variance to be σ_0^2) are $\hat{\mu}_0 = \bar{X}$ and $\hat{\sigma}_0^2 = \sigma_0^2$.

The Likelihood Ratio evaluates to:

$$\Lambda = \frac{(\sigma_0^2)^{-n/2} \exp\left(-\frac{1}{2\sigma_0^2} \sum (X_i - \bar{X})^2\right)}{(\hat{\sigma}^2)^{-n/2} e^{-n/2}} = \left(\frac{\hat{\sigma}^2}{\sigma_0^2}\right)^{n/2} \exp\left(-\frac{n\hat{\sigma}^2}{2\sigma_0^2} + \frac{n}{2}\right)$$

To make this recognizable, let's define a new variable $W = \frac{n\hat{\sigma}^2}{\sigma_0^2} = \frac{(n-1)S^2}{\sigma_0^2}$. We can rewrite the likelihood ratio entirely in terms of W:

$$\Lambda \propto W^{n/2} e^{-W/2}$$

Calculus shows that this function $\Lambda(W)$ looks like a hill; it peaks when $W = n$ and shrinks toward zero as W gets very small or very large. Thus, rejecting for small Λ is mathematically equivalent to rejecting when W is in the extreme tails.

Probabilistic Distribution & Rejection Region: Because we had to estimate the unknown mean μ using the sample mean $\bar{X}$, we lose one degree of freedom. Under H_0, our test statistic is distributed as:

$$W = \frac{(n-1)S^2}{\sigma_0^2} \sim \chi^2_{n-1}$$

We reject H_0 if the variance is wildly off in either direction: $W \leq \chi^2_{1-\alpha/2, n-1}$ or $W \geq \chi^2_{\alpha/2, n-1}$.

Problem 3: Testing the Parameter of a Poisson Distribution

Let $X_1, X_2, \ldots, X_n \sim \text{Poisson}(\lambda)$ be a random sample.

(a) **One-Sided Test:**

We test $H_0 : \lambda \leq \lambda_0$ against $H_1 : \lambda > \lambda_0$. First, we check a simple hypothesis λ_0 vs λ_1 (where $\lambda_1 > \lambda_0$). The NP Lemma gives the likelihood ratio:

$$\frac{L(\lambda_1)}{L(\lambda_0)} = \frac{e^{-n\lambda_1}\lambda_1^{\sum X_i} / \prod X_i!}{e^{-n\lambda_0}\lambda_0^{\sum X_i} / \prod X_i!} = e^{-n(\lambda_1-\lambda_0)}\left(\frac{\lambda_1}{\lambda_0}\right)^{\sum X_i} \geq k$$

Taking the natural log gives:

$$-n(\lambda_1 - \lambda_0) + \left(\sum_{i=1}^{n} X_i\right) \ln\left(\frac{\lambda_1}{\lambda_0}\right) \geq \ln k$$

Because $\lambda_1 > \lambda_0$, the logarithmic term $\ln(\lambda_1/\lambda_0)$ is strictly positive. This means the likelihood ratio is monotonically increasing with respect to the sum of our observations. Thus, we reject H_0 for large values of the sum.

Probabilistic Distribution & Rejection Region: A fundamental property of the Poisson distribution is that the sum of independent Poisson variables is also Poisson. Let our test statistic be $Y = \sum_{i=1}^{n} X_i$. Under H_0, the parameter is additive, so $Y \sim \text{Poisson}(n\lambda_0)$. We reject H_0 if $Y \geq c$, where c is the smallest integer threshold ensuring the Type I error $P(Y \geq c \mid n\lambda_0) \leq \alpha$.

(b) **Two-Sided Test:**

We test $H_0 : \lambda = \lambda_0$ against $H_1 : \lambda \neq \lambda_0$.

Probabilistic Distribution & Rejection Region: The test statistic is once again the sum $Y = \sum_{i=1}^{n} X_i \sim \text{Poisson}(n\lambda_0)$ under H_0.

Unlike continuous distributions, the Poisson distribution is discrete. This means we cannot simply split α perfectly to get an exact Type I error probability. Instead, we use the cumulative distribution function $F(y; \mu)$ to find conservative critical integer bounds c_1 and c_2 such that the error strictly stays under our α threshold:

$$P(Y \leq c_1 \mid n\lambda_0) \leq \alpha/2 \quad \text{and} \quad P(Y \geq c_2 \mid n\lambda_0) \leq \alpha/2$$

We reject H_0 if our observed sum falls into either tail: $Y \leq c_1$ or $Y \geq c_2$.

Problem 4: Difference Between Two Normal Means

Let $X_1, \ldots, X_{n_1} \sim N(\mu_1, \sigma^2)$ and $Y_1, \ldots, Y_{n_2} \sim N(\mu_2, \sigma^2)$ be two completely independent random samples drawn from populations with equal variance.

(a) **Equal and Known Variance:**

We test $H_0 : \mu_1 - \mu_2 = 0$ against $H_1 : \mu_1 - \mu_2 > 0$. The most intuitive estimator for the difference in population means is the difference in sample means: $D = \bar{X} - \bar{Y}$.

Under H_0, we expect $E[D] = 0$. What about the variance? Because the two samples are independent, the variance of their difference is simply the sum of their individual variances:

$$\text{Var}(D) = \text{Var}(\bar{X}) + \text{Var}(\bar{Y}) = \frac{\sigma^2}{n_1} + \frac{\sigma^2}{n_2} = \sigma^2\left(\frac{1}{n_1} + \frac{1}{n_2}\right)$$

Probabilistic Distribution & Rejection Region: By standardizing the difference, our test statistic follows a Standard Normal distribution under H_0:

$$Z = \frac{\bar{X} - \bar{Y}}{\sigma\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}} \sim N(0,1)$$

We reject H_0 if $Z \geq z_\alpha$.

(b) **Equal but Unknown Variance:**

We test $H_0 : \mu_1 - \mu_2 = 0$ vs $H_1 : \mu_1 - \mu_2 \neq 0$. When deriving the LRT for this scenario, the unrestricted MLEs are the individual sample means $\hat{\mu}_1 = \bar{X}$, $\hat{\mu}_2 = \bar{Y}$, and the overall pooled variance $\hat{\sigma}^2$. Under the null hypothesis, we force the means to be equal, creating a single common mean MLE: $\hat{\mu}_0 = \frac{n_1\bar{X}+n_2\bar{Y}}{n_1+n_2}$.

Through the sum of squares decomposition, the Likelihood Ratio reduces to a function of the differences between the sample means:

$$\Lambda = \left(\frac{1}{1 + \frac{n_1 n_2}{n_1+n_2} \frac{(\bar{X}-\bar{Y})^2}{\sum(X_i-\bar{X})^2+\sum(Y_i-\bar{Y})^2}} \right)^{(n_1+n_2)/2}$$

To make this practical, we define the pooled sample variance S_p^2. This is essentially a weighted average of the two individual sample variances. By combining both samples to estimate the single unknown variance σ^2, we gain more precision:

$$S_p^2 = \frac{\sum_{i=1}^{n_1}(X_i - \bar{X})^2 + \sum_{i=1}^{n_2}(Y_i - \bar{Y})^2}{n_1 + n_2 - 2}$$

Substituting S_p^2 into our Λ equation allows us to express the Likelihood Ratio strictly in terms of a classic T-statistic:

$$\Lambda = \left(1 + \frac{T^2}{n_1 + n_2 - 2}\right)^{-(n_1+n_2)/2} \quad \text{where} \quad T = \frac{\bar{X} - \bar{Y}}{S_p\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}$$

Rejecting for small values of Λ implies rejecting when $|T|$ is remarkably large.

Probabilistic Distribution & Rejection Region: Because we estimated two parameters (the two individual sample means) to calculate our pooled variance, we lose two degrees of freedom. Under H_0, our test statistic follows a t-distribution:

$$T \sim t_{n_1+n_2-2}$$

We reject H_0 if the difference is extreme in either direction: $|T| \geq t_{\alpha/2,\, n_1+n_2-2}$.