EE 503

Lecture Notes Part 13

Christopher Wayne Walker, Ph.D.

13.0 Statistics: Confidence Intervals

We would like to quantify the confidence we have that what we are trying to estimate lies in some interval about our estimate. Our introduction to this concept will involve bit error probability but this same analysis applies to any Bernoulli/binomial trials. Later in this section we will derive the usual normal (Gaussian) confidence interval that is applicable to more general cases.

Suppose we are trying to estimate the probability of a bit error. Let P_b be the true probability of bit error and let $\hat{P}_b$ be the estimated probability of bit error. Let X be the number of bit errors observed. Then

$$\hat{P}_b = \frac{X}{n}$$

where we have n total trials. Thus, X is a binomial random variable. Note that here a success is defined to be a bit error. We then have

$$E[X] = nP_b, \quad Var(X) = nP_b(1 - P_b).$$

So, provided $n\hat{P}_b \geq 5$ and $n(1 - \hat{P}_b) \geq 5$ we take

$$\frac{n\hat{P}_b - nP_b}{\sqrt{nP_b(1 - P_b)}} = \frac{\hat{P}_b - P_b}{\sqrt{\frac{P_b(1 - P_b)}{n}}} = Z$$

where $Z \sim N(0, 1)$ is standard normal by the central limit theorem. Since P_b is unknown we use $\hat{P}_b$ to get

$$\frac{\hat{P}_b - P_b}{\sqrt{\frac{\hat{P}_b(1 - \hat{P}_b)}{n}}} = Z.$$

We define $z_{\alpha/2}$ to denote the point on the z-axis such that the probability that Z is greater than this value is $\alpha/2$. Similarly, we define $-z_{\alpha/2}$ to denote the point on the z-axis such that the probability that Z is less than this value is $\alpha/2$. Hence, the probability of Z being in the interval $(-z_{\alpha/2}, z_{\alpha/2})$ is $1 - \alpha$. The points $\pm z_{\alpha/2}$ are called *critical values*.

To compute a confidence interval we then solve

$$-z_{\alpha/2} < \frac{\hat{P}_b - P_b}{\sqrt{\dfrac{\hat{P}_b(1-\hat{P}_b)}{n}}} < z_{\alpha/2}$$

to get a $100(1-\alpha)\%$ as

$$\hat{P}_b \pm z_{\alpha/2}\sqrt{\frac{\hat{P}_b(1-\hat{P}_b)}{n}} = \hat{P}_b \pm E$$

where

$$E = z_{\alpha/2}\sqrt{\frac{\hat{P}_b(1-\hat{P}_b)}{n}}$$

is referred to as the *margin of error*.

Using E we can solve for n to get

$$n = \frac{z_{\alpha/2}^2 \hat{P}_b(1-\hat{P}_b)}{E^2}.$$

In a political poll n would tell us how many people we need to poll and if we were going to count bit errors n would indicate how many bits we need to process or simulate. However, we need to know n before we conduct the experiment and $\hat{P}_b$ is not known at that point. Ideally, we would use

$$n = \frac{z_{\alpha/2}^2 P_b(1-P_b)}{E^2}$$

but, of course, P_b is not known either (this is what we are trying to estimate). However, $P_b(1-P_b)$ is maximized then $P_b = 0.5$ so we can conservatively use

$$n = \frac{z_{\alpha/2}^2 \cdot 0.25}{E^2}.$$

When counting things like bit errors it is often meaningful to consider E as a fraction of P_b. So if we replace E above with EP_b we get

$$n = \frac{z_{\alpha/2}^2 P_b(1-P_b)}{E^2 P_b^2} = \frac{z_{\alpha/2}^2 (1-P_b)}{E^2 P_b}.$$

But in counting bit errors $\hat{P}_b$ is usually very small so $(1 - P_b) \approx 1$. Hence, we take

$$n = \frac{z_{\alpha/2}^2}{E^2 P_b}$$

which we approximate using

$$n = \frac{z_{\alpha/2}^2}{E^2 \hat{P}_b}.$$

If we are then interested in how many errors to count, call this k, then $k = n\hat{P}_b$ yielding

$$k = \frac{z_{\alpha/2}^2}{E^2}.$$

Example: Suppose you want to determine how many errors to count to be 95% confidence that the true probability of bit error is with 10% of the observed probability of bit error. Then

$$k = \frac{(1.96)^2}{(0.1)^2} = 385.$$

Now we will derive a confidence interval applicable to a wide range of problems. Suppose $X \sim N(\mu, \sigma^2)$. Say for now that σ^2 is known but μ is known. Let

$$\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i.$$

Then $\bar{X}$ is an estimate of μ. We want

$$P(\bar{X} - a \leq \mu \leq \bar{X} + a) = 1 - P(|\bar{X} - \mu| > a) \geq 1 - \alpha.$$

The value of α is typically small, like 0.01 or 0.05. We find

$$P(|\bar{X} - \mu| > a) = 1 - P\left(\left|\frac{\sqrt{n}(\bar{X} - \mu)}{\sigma}\right| \leq \frac{a\sqrt{n}}{\sigma}\right).$$

But,

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

so

$$P(|\bar{X} - \mu| > a) = 1 - \Phi\left(\frac{a\sqrt{n}}{\sigma}\right) + \Phi\left(-\frac{a\sqrt{n}}{\sigma}\right) = 2\Phi\left(-\frac{a\sqrt{n}}{\sigma}\right) = \alpha$$

where $\Phi(z)$ is the standard normal cdf. So,

$$a = -\frac{\sigma}{\sqrt{n}}\Phi^{-1}(\alpha/2) = \frac{\sigma}{\sqrt{n}}\Phi^{-1}(1-\alpha/2) = \frac{\sigma}{\sqrt{n}}z_{\alpha/2}.$$

Hence,

$$I = \left(\bar{X} - \frac{\sigma}{\sqrt{n}}z_{\alpha/2}, \bar{X} + \frac{\sigma}{\sqrt{n}}z_{\alpha/2}\right)$$

is a level $(1-\alpha)$ confidence interval for μ. Note that in most situations if the mean is not known then the variance is not known either. In this case we have to use an estimate of the standard deviation in our confidence interval which leads to a t-distribution instead of standard normal. The t-distribution looks a lot like the standard normal but it depends on the sample size n and approaches the standard normal as $n \to \infty$. Here our confidence interval becomes

$$I = \left(\bar{X} - \frac{s}{\sqrt{n}}t_{\alpha/2}, \bar{X} + \frac{s}{\sqrt{n}}t_{\alpha/2}\right)$$

where s^2 is the usual unbiased estimator of σ^2.