Classical and Bayesian inference

AMS 132

February 26, 2018

Confidence Intervals for the Mean of a Normal Distribution

- Confidence intervals provide a method of adding more information to an estimator $\widehat{\theta}$ when we wish to estimate an unknown parameter θ. We can find an interval (A, B) that we think has high probability of containing θ. The length of such an interval gives us an idea of how closely we can estimate θ.

Example

Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean μ and variance σ^2. We know that an estimator for μ is $\overline{X}_n$, but how close is it from μ? How much confidence should we place in the estimator $\overline{X}_n$ for μ?

Confidence Intervals for the Mean of a Normal Distribution

a) Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean μ and variance σ^2.

b) Consider $\overline{X}_n$ and σ' estimators for μ and σ, respectively. What is the distribution of $U = \frac{\sqrt{n}(\overline{X}_n - \mu)}{\sigma'}$?

c) Find c such that $P(-c < U < c) = 0.95$. More generally, Find c such that $P(-c < U < c) = \gamma$, $0 < \gamma < 1$.

d) Use c) to find the random variables A and B, such that $P(A < \mu < B) = 0.95$.

Confidence Intervals for the Mean of a Normal Distribution

Definition (Confidence Interval)

Let $\boldsymbol{X} = (X_1, \ldots, X_n)$ be a random sample from a distribution that depends on a parameter (or parameter vector) θ. Let $g(\theta)$ be a real-valued function of θ. Let $A \leq B$ be two statistics that have the property that for all values of θ,

$$Pr(A < g(\theta) < B) \geq \gamma.$$

Then the random interval (A, B) is called a *coefficient* γ *confidence interval for* $g(\theta)$ or a 100γ *percent confidence interval for* $g(\theta)$.

If the inequality "$\geq \gamma$" is replaced by an equality for all θ, the confidence interval is called *exact*.

After the values of the random variables $X_1, \ldots, X_n$ in the random sample have been observed, the values of $A = a$ and $B = b$ are computed, and the interval (a, b) is called the *observed value of the confidence interval*.

Confidence Intervals for the Mean of a Normal Distribution

Theorem (Confidence Interval for the Mean of a Normal Distribution)

Let $X_1, \dots, X_n$ be a random sample from the normal distribution with mean μ and variance σ^2. For each $0 < \gamma < 1$, the interval (A, B) with the following endpoints is an exact coefficient γ confidence interval for μ:

$$A = \overline{X}_n - T_{n-1}^{-1}\left(\frac{1+\gamma}{2}\right)\frac{\sigma'}{\sqrt{n}},$$

$$B = \overline{X}_n + T_{n-1}^{-1}\left(\frac{1+\gamma}{2}\right)\frac{\sigma'}{\sqrt{n}}.$$

Confidence Intervals for the Mean of a Normal Distribution

Example (Lactic acid concentration in cheese)

One chemical whose concentration can affect taste is lactic acid. Cheese manufacturers who want to establish a loyal customer base would like the taste to be about the same each time a customer purchases the cheese. The variation in concentrations of chemicals like lactic acid can lead to variation in the taste of cheese.

Suppose that we model the concentration of lactic acid in several chunks of cheese as independent normal random variables with mean μ and variance σ^2, where μ and σ^2 are unknown.

a) Find a 90 percent confidence interval for μ, the unknown mean lactic acid concentration.
b) Suppose that we observe the following lactic acid concentrations $0.86, 1.53, 1.57, 1.81, 0.99, 1.09, 1.29, 1.78, 1.29, 1.58$. Compute the observed value of the 90 percent confidence interval for μ.
c) Make an interpretation of the intervals found in a) and b).

Confidence Intervals for the Mean of a Normal Distribution: interpretation

a) Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean $\mu = 5$ and variance $\sigma^2 = 2.5$.

b) Find a 90 percent confidence interval for the mean (we know the mean!).

c) Generate a sample of size 26 from the sampling distribution and compute the observed 90 percent confidence interval for the mean. Check if μ is in the computed interval or not.

d) Repeat c)

e) Repeat c) 100 times and each time check if μ is in the computed interval or not.

One-Sided Confidence Intervals

- There are situations where only an upper or lower bound for μ is of interest. Example: lactic acid concentration in cheese.
- To do so, we have to extend the definition of confidence interval to allow either $A = -\infty$ or $B = \infty$ so that the confidence interval has the form $(-\infty, B)$ or (A, ∞), respectively.

One-Sided Confidence Intervals

Definition (One-Sided Confidence Interval)

Let $\boldsymbol{X} = (X_1, \ldots, X_n)$ be a random sample from a distribution that depends on a parameter (or parameter vector) θ. Let $g(\theta)$ be a real-valued function of θ. Let A be a statistic that has the property that for all values of θ,

$$Pr(A < g(\theta)) \geq \gamma.$$

Then the random interval (A, ∞) is called a *one-sided* 100γ *percent confidence interval for* $g(\theta)$. Also, A is called a 100γ *percent lower confidence limit for* $g(\theta)$.

Similarly, if B is a statistic such that

$$Pr(g(\theta) < B) \geq \gamma,$$

then the random interval $(-\infty, B)$ is called a *one-sided* 100γ *percent confidence interval for* $g(\theta)$. Also, B is called a 100γ *percent upper confidence limit for* $g(\theta)$
If the inequality "$\geq \gamma$" is replaced by an equality for all θ, the confidence intervals are called *exact*.

One-Sided Confidence Intervals

- The interval (A, ∞) is also called *one-sided coefficient* γ *confidence interval for* $g(\theta)$, and A is also named *coefficient* γ *lower confidence limit for* $g(\theta)$.
- The interval $(-\infty, B)$ is also called *one-sided coefficient* γ *confidence interval for* $g(\theta)$, and B is also named *coefficient* γ *upper confidence limit for* $g(\theta)$.
- After the values of the random variables $X_1, \ldots, X_n$ in the random sample have been observed, the values of $A = a$ and $B = b$ are computed, and the intervals (a, ∞) and (∞, b) are called the *observed value of the confidence interval*.

One-Sided Confidence Intervals

Theorem (One-sided Confidence Interval for the Mean of a Normal Distribution)

Let $X_1, \ldots, X_n$ be a random sample from the normal distribution with mean μ and variance σ^2. For each $0 < \gamma < 1$, the following statistics are, respectively, exact lower and upper coefficient γ confidence limits for μ:

$$A = \overline{X}_n - T_{n-1}^{-1}(\gamma)\,\frac{\sigma'}{\sqrt{n}},$$

$$B = \overline{X}_n + T_{n-1}^{-1}(\gamma)\,\frac{\sigma'}{\sqrt{n}}.$$

One-Sided Confidence Intervals for the Mean of a Normal Distribution

Find a 100γ percent lower confidence limit for μ:

a) Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean μ and variance σ^2.
b) Consider $\overline{X}_n$ and σ' estimators for μ and σ, respectively. What is the distribution of $U = \frac{\sqrt{n}(\overline{X}_n - \mu)}{\sigma'}$? Justify each step.
c) Find c such that $P(U < c) = \gamma$, $0 < \gamma < 1$.
d) Use c) to find the random variable A, such that $P(A < \mu) = \gamma$.

Find a 100γ percent upper confidence limit for μ:

a) Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean μ and variance σ^2.
b) Consider $\overline{X}_n$ and σ' estimators for μ and σ, respectively. What is the distribution of $U = \frac{\sqrt{n}(\overline{X}_n - \mu)}{\sigma'}$? Justify each step.
c) Find c such that $P(U > -c) = \gamma$, $0 < \gamma < 1$.
d) Use c) to find the random variable B, such that $P(\mu < B) = \gamma$.

Confidence Intervals for Other Parameters

- Suppose that $X_1, \ldots, X_n$ form a random sample from the normal distribution with mean μ and variance σ^2. How can we find a confidence interval for the variance?

Definition (Pivotal)

Let $\boldsymbol{X} = (X_1, \ldots, X_n)$ be a random sample from a distribution that depends on a parameter (or vector of parameters) θ. Let $V(\boldsymbol{X}, \theta)$ be a random variable whose distribution is the same for all θ. Then V is called a *pivotal quantity* (or simply a *pivotal*).

- In order to use a pivotal for constructing a confidence interval for $g(\theta)$, we need to invert the pivotal. This is, wee need a function $r(v, \boldsymbol{x})$ such that

$$r(V(\boldsymbol{X}, \theta), \boldsymbol{X}) = g(\theta).$$

Confidence Intervals for Other Parameters

Theorem (Confidence interval from a Pivotal)

Let $\boldsymbol{X} = (X_1, \ldots, X_n)$ be a random sample from a distribution that depends on a parameter (or vector of parameters) θ. Suppose that a pivotal V exists. Let G be the c.d.f. of V, and assume that G is continuous. Assume that a function r exists as previously defined, and assume that $r(v, \boldsymbol{x})$ is strictly increasing in v for each $\boldsymbol{x}$. Let $0 < \gamma < 1$ and let $\gamma_2 > \gamma_1$ be such that $\gamma_2 - \gamma_1 = \gamma$. Then the following statistics are the endpoints of an exact coefficient γ confidence interval for $g(\theta)$:

$$A = r\left(G^{-1}(\gamma_1), \boldsymbol{X}\right),$$

$$B = r\left(G^{-1}(\gamma_2), \boldsymbol{X}\right).$$

If $r(v, \boldsymbol{x})$ is strictly decreasing in v for each $\boldsymbol{x}$, then switch the definitions of A and B.

Confidence Intervals for Other Parameters

Example

Suppose that $X_1, \ldots, X_{21}$ form a random sample from the normal distribution with mean μ and variance σ^2. Find a 90 percent confidence interval for the variance. For this:

a) Show that $Y = \sum_{i=1}^{n} \left(\frac{X_i - \overline{X}_n}{\sigma} \right)^2$ is a pivotal.

b) Consider $\gamma_1 = 0.05$ and $\gamma_2 = 0.95$, and compute $G^{-1}(\gamma_1)$ and $G^{-1}(\gamma_2)$, where G is the c.d.f. of Y.

c) Find random variables A and B such that $P(A < \sigma^2 < B) = \gamma$.

Confidence Intervals for Other Parameters

Example

Suppose that $X_1, \ldots, X_7$ describe the age of 7 male students in a class and $Y_1, \ldots, Y_7$ describe the age of 7 female students in a class. Assume that $X_i \overset{iid}{\sim} N(\mu_1, \sigma^2)$, $Y_i \overset{iid}{\sim} N(\mu_2, \sigma^2)$, they are independent, and the variance is known an equal to 2. It is of interest to find a 90 percent confidence interval for the difference of the mean age between males and females. For this:

a) Show that $Y = \frac{\sqrt{n}(\overline{X}_n - \overline{Y}_n - [\mu_1 - \mu_2])}{\sqrt{2}\sigma}$ is a pivotal.
b) Consider $\gamma_1 = 0.05$ and $\gamma_2 = 0.95$, and compute $G^{-1}(\gamma_1)$ and $G^{-1}(\gamma_2)$, where G is the c.d.f. of Y.
c) Find random variables A and B such that $P(A < \mu_1 - \mu_2 < B) = \gamma$.
d) If $\overline{x}_n = 18.49$ and $\overline{y}_n = 20.66$, compute the γ percent confidence interval. What can you say about ages of males and females in the class?