STAT 713 sp 2023 Lec 14 slides

Confidence intervals and sets

Karl B. Gregory

University of South Carolina

These slides are an instructional aid; their sole purpose is to display, during the lecture, definitions, plots, results, etc. which take too much time to write by hand on the blackboard. They are not intended to explain or expound on any material.

1. Confidence intervals and sets
2. Confidence intervals from pivotal quantities
3. Confidence sets from inverting a test of hypotheses
4. Optimality of confidence intervals
5. Bayesian credible intervals

Set and interval estimators

Let data $\mathbf{X}$ have a distribution depending on $\theta \in \Theta$ and consider estimating $\tau(\theta)$.

1. A set $C(\mathbf{X})$ intended to contain $\tau(\theta)$ is called a *set estimator* for $\tau(\theta)$.
2. $C(\mathbf{X}) = [L(\mathbf{X}), U(\mathbf{X})]$, then $C(\mathbf{X})$ is called an *interval estimator* for $\tau(\theta)$.

Typically $C(\mathbf{X}) \subset \mathbb{R}^d$ where $d = \dim(\{\tau(\theta) : \theta \in \Theta\})$.

An interval estimator is a special case of a set estimator.

Coverage probability and confidence level

Let $C = C(\mathbf{X})$ be a set estimator for $\tau(\theta)$.

1. The *coverage probability* of C under the value θ is $P_\theta(C(\mathbf{X}) \ni \tau(\theta))$.
2. The *confidence level* of C is $\inf_{\theta \in \Theta} P_\theta(C(\mathbf{X}) \ni \tau(\theta))$.

The confidence level is the smallest probability of coverage over all $\theta \in \Theta$.

A set estimator with conf. level $\geq 1 - \alpha$ is called a $(1-\alpha) \times 100\%$ *confidence set*.

We will discuss two ways of constructing confidence sets/intervals:

1. Use a *pivotal quantity*.
2. Invert a test of hypotheses.

Pivotal quantity

A rv $Q(\mathbf{X};\theta)$ is a *pivotal quantity* for θ if $P_\theta(Q(\mathbf{X};\theta) \in \mathcal{A})$ is free of θ for all $\mathcal{A}$.

I.e. $Q(\mathbf{X};\theta)$ is a *pivotal quantity* for θ if its distribution does not depend on θ.

Easy to find pivotal quantities for location/scale parameters (pg. 427 of CB).

Theorem (How to use a pivotal quantity to build a confidence set)

Let $Q(\mathbf{X};\theta)$ be pivotal for θ and $\mathcal{A}$ a set st $P_\theta(Q(\mathbf{X};\theta) \in \mathcal{A}) \geq 1-\alpha$. Then

$$C(\mathbf{X}) = \{\theta : Q(\mathbf{X};\theta) \in \mathcal{A}\}$$

is a $(1-\alpha)\times 100\%$ confidence set for θ.

Exercise: Prove the result.

Exercise: Let $X_1, \dots, X_n \stackrel{\text{ind}}{\sim} \text{Exponential}(\lambda)$.

1. Find a pivotal quantity for λ.
2. Find a $(1 - \alpha) \times 100\%$ confidence set for λ.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$. Build confidence sets for (μ, σ^2) with

1. The pivot $Q(\mathbf{X}; \mu, \sigma^2) = \sqrt{n}(\bar{X}_n - \mu)/\sigma$ and the set $\mathcal{A} = [-z_{\alpha/2}, z_{\alpha/2}]$.
2. The pivot $Q(\mathbf{X}; \mu, \sigma^2) = \left(\sqrt{n}(\bar{X}_n - \mu)/S_n, (n-1)S_n^2/\sigma^2\right)$ and the set

$$\mathcal{A} = [-t_{n-1,\alpha/2}, t_{n-1,\alpha/2}] \times [\chi^2_{n-1,1-\alpha/2}, \chi^2_{n-1,\alpha/2}].$$

Check the confidence level.

Theorem (Pivoting a cdf)

Let $\theta \in \mathbb{R}$ and $T = T(\mathbf{X}) \in \mathbb{R}$ be continuous with cdf $F_T(t;\theta)$. If for each t the cdf $F_T(t;\theta)$ is continuous and monotone in θ then $Q(T;\theta) = F_T(T;\theta)$ is a pivotal quantity for θ. Moreover, if $\alpha_1 + \alpha_2 = \alpha$, then

$$\{\theta : \alpha_1 \leq F_T(T;\theta) \leq 1 - \alpha_2\}$$

is a $(1-\alpha) \times 100\%$ confidence interval for θ.

When T is discrete, the above may give a CI of level approximately $1-\alpha$.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Uniform}(0,\theta)$. Find a $(1-\alpha) \times 100\%$ for θ by pivoting the cdf of $X_{(n)}$.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Gamma}(\gamma, 1)$.

1. Find a $(1 - \alpha) \times 100\%$ for γ by pivoting the cdf of $T = \sum_{i=1}^{n} X_i$.
2. Do you think the interval is the best possible?

Theorem (Finding a confidence set by test inversion)

For each $\theta_0 \in \Theta$, let ϕ_{θ_0} be a level-α test of H_0: $\theta = \theta_0$. Then

$$C(\mathbf{X}) = \{\theta_0 : \phi_{\theta_0}(\mathbf{X}) = 0\}$$

is a $(1-\alpha) \times 100\%$ confidence set for θ.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$. Find a CI for μ by inverting

1. the test of H_0: $\mu = \mu_0$ vs H_1: $\mu \neq \mu_0$ with rejection rule $|T| > t_{n-1,\alpha/2}$.
2. the test of H_0: $\mu \leq \mu_0$ vs H_1: $\mu > \mu_0$ with rejection rule $T > t_{n-1,\alpha}$.

In the above $T = \sqrt{n}(\bar{X}_n - \mu_0)/S_n$.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} f(x; \beta) = \beta x^{-(\beta+1)} \mathbf{1}(x > 1)$. Invert the size-$\alpha$ LRT of H_0: $\beta = \beta_0$ vs H_1: $\beta \neq \beta_0$ to obtain a $(1 - \alpha) \times 100\%$ confidence interval for β.

Familywise coverage probability

Given K confidence sets, let C_k be the event that confidence set k covers its target for $k = 1, \ldots, K$. The probability $P(\cap_{k=1}^K C_k)$ is called the *familywise coverage probability* for the K confidence sets.

Bonferroni method for ensuring a familywise coverage probability

If C_k is the coverage event for CI k and $P(C_k) = 1 - \alpha$ for all $k = 1, \ldots, K$, then

$$P(\cap_{k=1}^K C_k) \geq 1 - K\alpha.$$

So if $\alpha = \alpha^*/K$ then the familywise coverage probability will be at least $1 - \alpha^*$.

Exercise: Derive the Bonferroni method.

Exercise: Let $X_{k1}, \dots, X_{kn_k} \overset{\text{ind}}{\sim} \text{Normal}(\mu_k, \sigma^2)$ be two indep. random samples. Give CIs for μ_1, μ_2, and $\mu_1 - \mu_2$ with familywise coverage at least $1 - \alpha$.

Asymptotic confidence sets

Let $C_n(\mathbf{X}_n)$ be a sequence of set estimators of $\tau(\theta)$. If

$$\inf_{\theta\in\Theta}\lim_{n\to\infty} P_\theta(C_n(\mathbf{X}_n)\ni\tau(\theta))\geq 1-\alpha,$$

we call $C_n(\mathbf{X}_n)$ an *asymptotic* $(1-\alpha)\times 100\%$ *confidence set* for $\tau(\theta)$.

Asymptotic pivotal quantitity

Let $Q_n(\mathbf{X}_n;\theta)$ be a seq. of rvs such that $Q_n(\mathbf{X}_n;\theta)\xrightarrow{\text{D}} Q$ as $n\to\infty$, where the distribution of Q is free of θ. Then $Q_n(\mathbf{X}_n;\theta)$ is an *asymptotic pivotal quantity*.

In this case, if $\mathcal{A}$ is a set such that $P(Q\in\mathcal{A})\geq 1-\alpha$, then

$$C_n(\mathbf{X}_n)=\{\theta : Q_n(\mathbf{X}_n;\theta)\in\mathcal{A}\}$$

is an asymptotic $(1-\alpha)\times 100\%$ confidence set for θ.

Exercise: Let $X_1, \dots, X_n$ be iid with moment $m_{2j} < \infty$. Use the fact

$$\sqrt{n}(\hat{m}_j - m_j)/\sqrt{\hat{m}_{2j} - \hat{m}_j^2} \xrightarrow{\text{D}} Z \sim \text{Normal}(0, 1)$$

as $n \to \infty$ to build an asymptotic confidence interval for m_j.

Some likelihood-based confidence interval recipes

Let $\tau(\theta)$ be real-valued and assume the standard ML regularity conditions.

- If $\sqrt{n}(\tau(\hat{\theta}_n) - \tau(\theta)) \xrightarrow{\text{D}} \text{Normal}(0, \vartheta(\theta))$, where $\hat{\theta}_n$ is the MLE, then

  $$\tau(\hat{\theta}_n) \pm z_{\alpha/2}\sqrt{\vartheta(\hat{\theta}_n)/n}.$$

  is an asymptotic $(1-\alpha)\times 100\%$ CI for $\tau(\theta)$. Called a *Wald-type* CI.

- Let $W_n(\mathbf{X}_n; \tau_0)$ be the score or ALRT statistic for H_0: $\tau = \tau_0$ and suppose $W_n(\mathbf{X}_n; \tau_0) \xrightarrow{\text{D}} \chi^2_1$ under H_0. Then the interval defined by

  $$\{\tau_0 : W_n(\mathbf{X}_n; \tau_0) \leq \chi^2_{1,\alpha}\}$$

  is an asymptotic $(1-\alpha)\times 100\%$ CI for $\tau(\theta)$. Called a *score or ALRT* interval.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Poisson}(\lambda)$. For $\tau(\lambda) = 1/\lambda$, show:

1. The $1-\alpha$ Wald interval is $\dfrac{1}{\hat{\lambda}_n} \pm z_{\alpha/2}\sqrt{\dfrac{1}{n\hat{\lambda}_n^3}}$.
2. The $1-\alpha$ Score interval is $\dfrac{1}{\hat{\lambda}_n} + \dfrac{\chi^2_{1,\alpha}}{2n\hat{\lambda}_n^2} \pm z_{\alpha/2}\sqrt{\dfrac{1}{n\hat{\lambda}_n^3} + \dfrac{\chi^2_{1,\alpha}}{n^2\hat{\lambda}_n^4}}$.
3. The $1-\alpha$ ALR interval is $\left\{\tau : \dfrac{2n}{\tau}\left[(1-\tau\hat{\lambda}_n) + \tau\hat{\lambda}_n \log(\tau\hat{\lambda}_n) \leq \chi^2_{1,\alpha}\right]\right\}$.

In the above, $\hat{\lambda}_n$ is the MLE of λ.

Exercise: Let $X_{k1}, \ldots, X_{kn_k} \overset{\text{ind}}{\sim} \text{Poisson}(\lambda_k)$ for $k = 1, 2$ be independent samples. Show that the $1 - \alpha$ score interval for λ_1/λ_2 is given by

$$\frac{\hat{\lambda}_1}{\hat{\lambda}_2} + \left(\frac{n_1\hat{\lambda}_1 + n_2\hat{\lambda}_2}{2n_1 n_2 \lambda_2^2}\right)\chi^2_{1,\alpha} \pm z_{\alpha/2}\sqrt{\frac{\hat{\lambda}_1}{\hat{\lambda}_2}\left(\frac{n_1\hat{\lambda}_1 + n_2\hat{\lambda}_2}{n_1 n_2 \hat{\lambda}_2^2}\right) + \left(\frac{n_1\hat{\lambda}_1 + n_2\hat{\lambda}_2}{2n_1 n_2 \hat{\lambda}_2^2}\right)^2 \chi^2_{1,\alpha}}$$

False coverage probability and unbiasedness of confidence sets

Let $C(\mathbf{X})$ be a confidence set for $\tau(\theta)$, $\theta \in \Theta$.

- The *confidence level* of $C(\mathbf{X})$ is $\inf_{\theta \in \Theta} P_\theta(C(\mathbf{X}) \ni \tau(\theta))$.
- The *false coverage prob.* of $C(\mathbf{X})$ is $P_\theta(C(\mathbf{X}) \ni \tau(\theta'))$ for each $\theta, \theta' \in \Theta$.
- The set $C(\mathbf{X})$ is an *unbiased* if

$$\sup_{\theta,\theta' \in \Theta} P_\theta(C(\mathbf{X}) \ni \tau(\theta')) \leq \inf_{\theta \in \Theta} P_\theta(C(\mathbf{X}) \ni \tau(\theta)).$$

Unbiased $C(\mathbf{X})$ covers the true value more often than it covers any other value.

UMAU confidence sets

Let $C(\mathbf{X})$ be an unbiased $1-\alpha$ confidence set for τ. We call $C(\mathbf{X})$ *uniformly most accurate unbiased (UMAU)* if for any other unbiased $1-\alpha$ conf. set $\tilde{C}(\mathbf{X})$ we have

$$P_\theta(C(\mathbf{X}) \ni \tau(\theta')) \leq P_\theta(\tilde{C}(\mathbf{X}) \ni \tau(\theta')) \text{ for all } \theta, \theta' \in \Theta.$$

So $C(\mathbf{X})$ is UMAU if its false coverage probability is as small as possible.

Theorem (Invert UMPU test to obtain UMAU confidence set)

Suppose $\phi(\mathbf{X}) = \mathbf{1}(\tau_0 \notin C(\mathbf{X}))$ *is a size* α *UMPU test for* H_0*:* $\tau(\theta) = \tau_0$ *vs* H_1*:* $\tau(\theta) \neq \tau_0$*. Then* $C(\mathbf{X})$ *is a UMAU* $1-\alpha$ *confidence set for* $\tau(\theta)$*.*

Example: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$. UMPU size α test of H_0: $\mu = \mu_0$ vs H_1: $\mu \neq \mu_0$ is $\phi(X) = \mathbf{1}(|\sqrt{n}(\bar{X}_n - \mu_0)/S_n| > t_{n-1,\alpha/2})$. Inverting this gives UMAU $1-\alpha$ confidence set $[\bar{X}_n \pm t_{n-1,\alpha/2} S_n/\sqrt{n}]$.

Consider constructing a Bayesian set/interval estimator for θ in the model

$$\begin{aligned} \mathbf{X}|\theta &\sim f(\mathbf{x}|\theta) \\ \theta &\sim \pi(\theta), \quad \theta \in \Theta. \end{aligned}$$

Bayesian credible interval

A $(1-\alpha) \times 100\%$ *credible set* $C(\mathbf{X})$ is a set $C(\mathbf{X}) \subset \Theta$ such that

$$P(C(\mathbf{X}) \ni \theta|\mathbf{X}) = 1 - \alpha.$$

Theorem (Coverage probability of Bayesian credible interval)

If $C(\mathbf{X})$ is a $(1-\alpha) \times 100\%$ credible set, then its expected coverage probability over all values of $\theta \in \Theta$ is $1-\alpha$. That is $\mathbb{E}[P(C(\mathbf{X}) \ni \theta|\theta)] = 1-\alpha$.

Exercise: Prove the result.

Exercise: Consider the hierarchical model

$$\begin{aligned} Y_1, \dots, Y_n | \mu &\overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2) \\ \mu &\sim \text{Normal}(\mu_0, \tau^2), \end{aligned}$$

where σ^2 is known and μ_0 and τ^2 are prior parameters.

1. Find the posterior distribution of $\mu | Y_1, \dots, Y_n$.
2. Find a $(1 - \alpha) \times 100\%$ credible set for μ based on $Y_1, \dots, Y_n$.
3. Make pictures of the credible interval under $\sigma^2 = 10^2$, $\mu_0 = 100$, $\tau^2 = 20^2$, and $\bar{Y}_n = 120$ with the sample sizes $n = 1, 5, 15$.

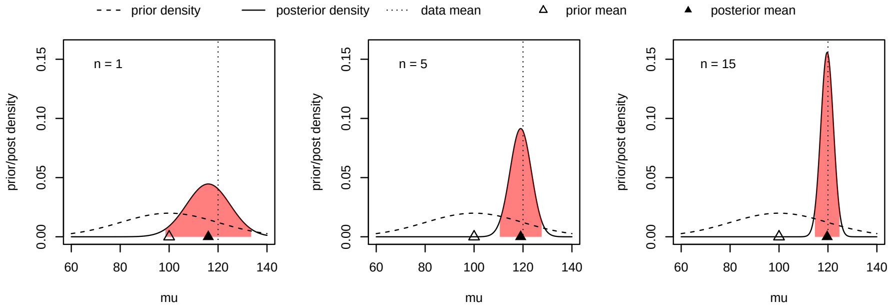
prior density
posterior density
data mean
prior mean
posterior mean
n = 1
n = 5
n = 15
prior/post density
0.00
0.05
0.10
0.15
60
80
100
120
140
mu

Highest posterior density credible set

The *highest posterior density (HPD)* $(1-\alpha)\times 100\%$ *credible set* for θ is the set $C(\mathbf{X})$ such that

1. $P(C(\mathbf{X}) \ni \theta|\mathbf{X}) = 1-\alpha$
2. $\pi(\theta|\mathbf{X}) > k \iff \theta \in C(\mathbf{X})$ for some $k>0$.

If $\theta \in \mathbb{R}$ and $\pi(\theta|\mathbf{X})$ unimodal, the HPD credible set is given by $[c_1, c_2]$, where

$$\int_{c_1}^{c_2} \pi(\theta|\mathbf{X})d\theta = 1-\alpha \quad \text{and} \quad \pi(c_1|\mathbf{X}) = \pi(c_2|\mathbf{X}).$$

Exercise: Suppose

$$X_1, \dots, X_n | \lambda \overset{\text{ind}}{\sim} \text{Poisson}(\lambda)$$
$$\lambda \sim \text{Gamma}(\alpha, \beta),$$

where α and β are prior parameters.

1. Find the posterior distribution of $\lambda | X_1, \dots, X_n$.
2. Find a $(1-\alpha) \times 100\%$ credible interval for λ.
3. Discuss finding the $(1-\alpha) \times 100\%$ HPD credible interval for λ.

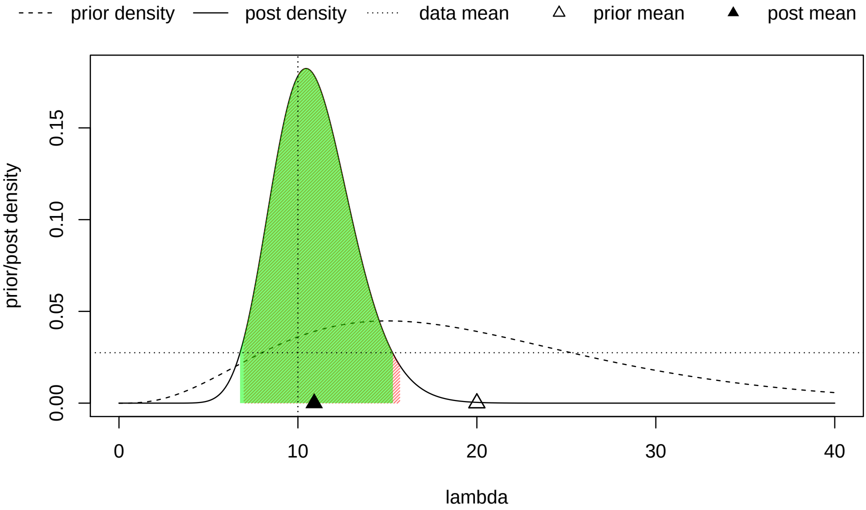
prior density
post density
data mean
prior mean
post mean
prior/post density
0.00
0.05
0.10
0.15
0
10
20
30
40
lambda