STK4011 and STK9011
Autumn 2016

Confidence intervals (Interval estimation)

Covers (most of) the following material from chapters 9 and 10:

- Section 9.1
- Section 9.2.1 (except examples 9.2.4 and 9.2.5)
- Section 10.4.1

Ørnulf Borgan
Department of Mathematics
University of Oslo

Basic concepts

Assume that we have random variables $\mathbf{X}=(X_1,X_2,\ldots,X_n)$ with joint pmf or pdf $f(\mathbf{x}\,|\,\theta)=f(x_1,\ldots,x_n\,|\,\theta)$ where $\theta\in\Theta$

We want to find a set $C\subset\Theta$ that contains θ

Definition 9.1.1

An interval estimate (or confidence interval) of a real valued parameter θ is a pair of functions, $L(\mathbf{x})$ and $U(\mathbf{x})$, that satisfy $L(\mathbf{x})\le U(\mathbf{x})$ for all $\mathbf{x}\in\mathcal{X}$. If $\mathbf{X}=\mathbf{x}$ is observed, we make the inference $L(\mathbf{x})\le\theta\le U(\mathbf{x})$. The random interval $[L(\mathbf{X})\,,U(\mathbf{X})]$ is called an interval estimator.

Definition 9.1.4

For an interval estimator $[L(\mathbf{X})\,,U(\mathbf{X})]$ of a parameter θ, the coverage probability is the probability that $[L(\mathbf{X})\,,U(\mathbf{X})]$ contains the true parameter θ, i.e. $P_\theta\big(\theta\in[L(\mathbf{X})\,,U(\mathbf{X})]\big)$

Definition 9.1.5

For an interval estimator $[L(\mathbf{X})\,,U(\mathbf{X})]$ of a parameter θ, the confidence coefficient is the infimum of the coverage probabilities, i.e. $\inf_{\theta\in\Theta} P_\theta\big(\theta\in[L(\mathbf{X})\,,U(\mathbf{X})]\big)$

A confidence interval with the confidence coefficient equal to $1-\alpha$ is called a $1-\alpha$ confidence interval

Example 9.1.6 (modified)

Let $X_1,\ldots,X_n$ be iid $\text{uniform}(\theta)$ where $\theta\in\Theta=[0.5,1.5]$

$Y=X_{(n)}=\max\{X_1,\ldots,X_n\}$ is a sufficient statistic, and it has pdf $f_Y(y)=ny^{n-1}/\theta^n \quad (0\le y\le\theta)$

Note that $T=Y/\theta$ has pdf $f_T(t)=nt^{n-1} \quad (0\le t\le 1)$

We will consider interval estimators of the form $[aY,bY]$ and $[Y+c,Y+d]$

Now we have

$$P_\theta\left(\theta\in[aY,bY]\right)=P_\theta\left(aY\le\theta\le bY\right)=P_\theta\left(\frac{1}{b}\le\frac{Y}{\theta}\le\frac{1}{a}\right)$$

$$=P_\theta\left(\frac{1}{b}\le T\le\frac{1}{a}\right)\;=\int_{1/b}^{1/a}nt^{n-1}dt\;=\left(\frac{1}{a}\right)^n-\left(\frac{1}{b}\right)^n$$

We choose a and b such that $(1/a)^n-(1/b)^n=0.95$

E.g. we may find a and b by solving

$$(1/a)^n=0.975 \quad \text{and} \quad (1/b)^n=0.025$$

Then the interval has coverage probability 95% for all $\theta\in\Theta=[0.5,1.5]$

Hence the confidence coefficient is also 95%

For the other interval we have

$$P_\theta\left(\theta\in[Y+c,Y+d]\right)=P_\theta\left(Y+c\le\theta\le Y+d\right)$$

$$=P_\theta\left(1-\frac{d}{\theta}\le T\le 1-\frac{c}{\theta}\right)\;=\left(1-\frac{c}{\theta}\right)^n-\left(1-\frac{d}{\theta}\right)^n$$

We may (e.g.) choose c and d by solving

$$(1-c)^n=0.975 \quad \text{and} \quad (1-d)^n=0.025$$

The figure shows the coverage probability as a function of θ for n = 10

The confidence coefficient is 88.3 %

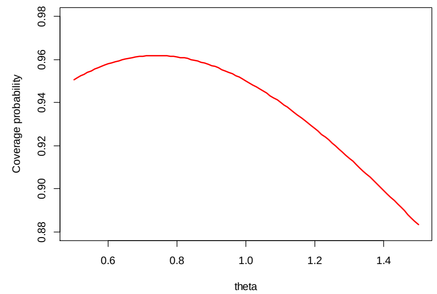

Confidence intervals and tests

There is a close connection between tests and confidence intervals (or more generally confidence sets)

Example 9.2.1 (inverting a normal test)

Let $X_1, X_2,....,X_n$ be iid $n(\mu,\sigma^2)$ with σ^2 known

Consider testing $H_0:\mu=\mu_0$ versus $H_1:\mu\ne\mu_0$

The most powerful unbiased level α test has rejection region

$$R(\mu_0)=\left\{\mathbf{x}:|\,\bar{x}-\mu_0\,|>z_{\alpha/2}\sigma/\sqrt{n}\right\}$$

The acceptance region is

$$A(\mu_0)=R(\mu_0)^c=\left\{\mathbf{x}:|\,\bar{x}-\mu_0\,|\le z_{\alpha/2}\sigma/\sqrt{n}\right\}$$

The test has size α so we have

$$P_{\mu_0}\left(\mathbf{X}\in A(\mu_0)\right)=1-P_{\mu_0}\left(\mathbf{X}\in R(\mu_0)\right)=1-\alpha$$

Now we have

$$\mathbf{X}\in A(\mu_0)\Leftrightarrow \bar{X}-z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\le\mu_0\le\bar{X}+z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

Thus

$$P_{\mu_0}\left(\bar{X}-z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\le\mu_0\le\bar{X}+z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\right)=1-\alpha$$

Since this is valid for all μ_0 we obtain

$$P_{\mu}\left(\bar{X}-z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\le\mu\le\bar{X}+z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\right)=1-\alpha$$

This gives the confidence interval

$$\left[\bar{x}-z_{\alpha/2}\sigma/\sqrt{n}\;,\bar{x}+z_{\alpha/2}\sigma/\sqrt{n}\right]$$

Theorem 9.2.2

For each $\theta_0 \in \Theta$, let $A(\theta_0) \subset \mathcal{X}$ be the acceptance region of a level α test of $H_0 : \theta = \theta_0$

For each $\mathbf{x} \in \mathcal{X}$, define $C(\mathbf{x}) \subset \Theta$ by

$$C(\mathbf{x}) = \{\theta_0 : \mathbf{x} \in A(\theta_0)\}$$

Then $C(\mathbf{X})$ is a confidence set with confidence coefficient $1-\alpha$ (i.e. a $1-\alpha$ confidence set)

Conversely, let $C(\mathbf{X})$ be a $1-\alpha$ confidence set, and for any $\theta_0 \in \Theta$, define

$$A(\theta_0) = \{\mathbf{x} : \theta_0 \in C(\mathbf{x})\}$$

Then $A(\theta_0)$ is the acceptance region for a level α test of $H_0 : \theta = \theta_0$

If we in theorem 9.2.2 consider a two-sided test, we will usually obtain a confidence interval (but there is no guarantee), while a one-sided test usually will give a one-sided confidence interval (i.e. the lower limit is $-\infty$ or the upper limit is ∞)

Example 9.2.3 (inverting the LRT)

Let $X_1, X_2, ..., X_n$ be iid and $\text{exponential}(\lambda)$

We want a $1-\alpha$ confidence interval of the mean λ

We can obtain such an interval by inverting a level α test of $H_0 : \lambda = \lambda_0$ versus $H_1 : \lambda \neq \lambda_0$

We will here use the LRT

The unrestricted MLE of λ is $\hat{\lambda} = \bar{x}$

The LRT statistic becomes

$$\frac{L(\lambda_0 \mid \mathbf{x})}{L(\hat{\lambda} \mid \mathbf{x})} = \frac{\prod_{i=1}^n (1/\lambda_0) e^{-x_i/\lambda_0}}{\prod_{i=1}^n (1/\hat{\lambda}) e^{-x_i/\hat{\lambda}}} = \left(\frac{\sum_{i=1}^n x_i}{n\lambda_0} \right)^n e^n e^{-\sum x_i/\lambda_0}$$

For fixed λ_0 the acceptance region is

$$A(\lambda_0) = \left\{ \mathbf{x} : \left(\sum_{i=1}^n x_i \Big/ \lambda_0 \right)^n e^{-\sum x_i/\lambda_0} \geq k^* \right\}$$

k^* is chosen such that

$$P_{\lambda_0}\left(\mathbf{X} \in A(\lambda_0)\right) = 1-\alpha$$

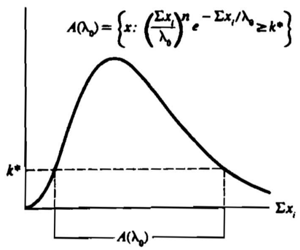

From this we obtain the confidence set

$$C(\mathbf{x}) = \left\{ \lambda : \left(\sum_{i=1}^n x_i \Big/ \lambda \right)^n e^{-\sum x_i/\lambda} \geq k^* \right\}$$

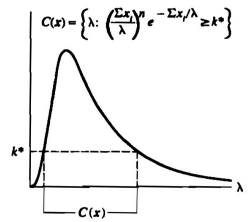

Note that $\lambda \in C(\mathbf{x})$ if and only if $L \leq \lambda \leq U$, where $L < U$ satisfy

$$\left(\sum_{i=1}^n x_i \Big/ L \right)^n e^{-\sum x_i/L} = \left(\sum_{i=1}^n x_i \Big/ U \right)^n e^{-\sum x_i/U}$$

We set $\sum_{i=1}^n x_i \Big/ L = a$ and $\sum_{i=1}^n x_i \Big/ U = b$

Then the confidence interval takes the form

$$\left[\frac{1}{a} \sum_{i=1}^n x_i \, , \, \frac{1}{b} \sum_{i=1}^n x_i \right]$$

where $b < a$ should satisfy $a^n e^{-a} = b^n e^{-b}$

Further a and b should be chosen such that the confidence coefficient becomes $1-\alpha$

Now $\sum_{i=1}^{n} X_i \sim \text{gamma}(n,\lambda)$ and $\sum_{i=1}^{n} X_i/\lambda \sim \text{gamma}(n,1)$

If we let $G_n(t)$ denote the cumulative $\text{gamma}(n,1)$ distribution, we have

$$P_\lambda\left(\frac{1}{a}\sum_{i=1}^n X_i \le \lambda \le \frac{1}{b}\sum_{i=1}^n X_i\right) = P_\lambda\left(b \le \sum_{i=1}^n X_i/\lambda \le a\right)$$

$$= G_n(a) - G_n(b)$$

So $b < a$ are the solutions of the equations

$$a^n e^{-a} = b^n e^{-b} \quad \text{and} \quad G_n(a) - G_n(b) = 1-\alpha$$

Approximate maximum likelihood intervals

Let $X_1,\ldots,X_n$ iid random variables with pdf or pmf $f(x|\theta)$ and let $\hat\theta$ be the MLE of θ

We may estimate the variance of a function $h(\hat\theta)$ using expected information

$$\widehat{\text{Var}}\, h(\hat\theta) = \frac{\left[h'(\theta)\right]^2\big|_{\theta=\hat\theta}}{nI_1(\theta)\big|_{\theta=\hat\theta}}$$

or using observed information

$$\widehat{\text{Var}}\, h(\hat\theta) = \frac{\left[h'(\theta)\right]^2\big|_{\theta=\hat\theta}}{-\sum_{i=1}^n \left(\partial^2/\partial\theta^2\right)\log f(X_i|\theta)\big|_{\theta=\hat\theta}}$$

(cf. slide 12 of week 43)

In any case we have (under some regularity conditions)

$$\frac{h(\hat\theta)-h(\theta)}{\sqrt{\widehat{\text{Var}}\, h(\hat\theta)}} \to Z \sim n(0,1)$$

It follows that

$$P_\theta\left(\left|\frac{h(\hat\theta)-h(\theta)}{\sqrt{\widehat{\text{Var}}\, h(\hat\theta)}}\right| \le z_{\alpha/2}\right) \to P\left(|Z| \le z_{\alpha/2}\right) = 1-\alpha$$

Hence

$$h(\hat\theta) - z_{\alpha/2}\sqrt{\widehat{\text{Var}}\, h(\hat\theta)} \le h(\theta) \le h(\hat\theta) + z_{\alpha/2}\sqrt{\widehat{\text{Var}}\, h(\hat\theta)}$$

is asymptotically a $1-\alpha$ confidence interval for $h(\theta)$

Example 10.4.1 (interval for the odds)

$X_1,\ldots,X_n$ are iid Bernoulli random variables with success probability p

ML estimator $\hat p = \sum_{i=1}^n X_i/n$

Consider estimation of the odds $h(p) = p/(1-p)$

Both methods for estimating the variance give

$$\widehat{\text{Var}}\left(\frac{\hat p}{1-\hat p}\right) = \frac{\hat p}{n(1-\hat p)^3}$$

An approximate $1-\alpha$ confidence interval for the odds is given by

$$\frac{\hat p}{1-\hat p} - z_{\alpha/2}\sqrt{\frac{\hat p}{n(1-\hat p)^3}} \le \frac{p}{1-p} \le \frac{\hat p}{1-\hat p} + z_{\alpha/2}\sqrt{\frac{\hat p}{n(1-\hat p)^3}}$$

One may often obtain good confidence intervals by using the score statistic

$$Q(\mathbf{X}\,|\,\theta) = \frac{\frac{\partial}{\partial\theta}\log L(\theta\,|\,\mathbf{X})}{\sqrt{-\mathrm{E}_\theta\left(\frac{\partial^2}{\partial\theta^2}\log L(\theta\,|\,\mathbf{X})\right)}} = \frac{\sum_{i=1}^{n}\frac{\partial}{\partial\theta}\log f(X_i\,|\,\theta)}{\sqrt{n\,I_1(\theta)}}$$

which has a $n(0,1)$ distribution asymptotically

By inverting the score test, we find that

$$\left\{\theta : \left|Q(\mathbf{X}\,|\,\theta)\right| \le z_{\alpha/2}\right\}$$

is an approximate $1-\alpha$ confidence set for θ

Example 10.4.2 (binomial score interval)

$X_1,\ldots,X_n$ are iid Bernoulli random variables with success probability p

ML estimator $\hat{p} = \sum_{i=1}^{n} X_i / n$

The score statistic is here (cf. example 10.3.5)

$$Q(\mathbf{X}\,|\,p) = \frac{\hat{p}-p}{\sqrt{p(1-p)/n}}$$

An approximate $1-\alpha$ confidence interval for *p* is given by

$$\left\{p : \left|\frac{\hat{p}-p}{\sqrt{p(1-p)/n}}\right| \le z_{\alpha/2}\right\}$$

See example 10.4.6 for an explicit expression

Another possibility is to use the LRT to obtain a confidence interval. We then use that

$$-2\log\lambda(\mathbf{X}) = -2\log\left(\frac{L(\theta\,|\,\mathbf{X})}{L(\hat{\theta}\,|\,\mathbf{X})}\right) \to \chi_1^2$$

Thus the set

$$\left\{\theta : -2\log\left(\frac{L(\theta\,|\,\mathbf{X})}{L(\hat{\theta}\,|\,\mathbf{X})}\right) \le \chi_{1,\alpha}^2\right\}$$

is an approximate $1-\alpha$ confidence set for θ

Example 10.4.3 (binomial LRT interval)

$X_1,\ldots,X_n$ are iid Bernoulli random variables with success probability p

ML estimator $\hat{p} = \sum_{i=1}^{n} X_i / n$

The LRT statistic is here

$$-2\log\lambda(\mathbf{X}) = -2\log\left(\frac{p^{\sum X_i}(1-p)^{n-\sum X_i}}{\hat{p}^{\sum X_i}(1-\hat{p})^{n-\sum X_i}}\right)$$

An approximate $1-\alpha$ confidence interval for *p* is given by

$$\left\{p : -2\log\left(\frac{p^{\sum X_i}(1-p)^{n-\sum X_i}}{\hat{p}^{\sum X_i}(1-\hat{p})^{n-\sum X_i}}\right) \le \chi_{1,\alpha}^2\right\}$$