Parametric Statistics

t Distributions, Confidence Intervals

Sofia Triantafillou

sof.triantafillou@gmail.com

University of Crete
Department of Mathematics and Applied Mathematics

November 27, 2023

Lecture Summary

Last time

- Let $X_1, \dots, X_n$ be a random sample from a $\mathcal{N}(\mu, \sigma^2)$ with unknown μ, σ^2.
- The sample mean and the sample variance are defined as

$$\overline{X}_n = \frac{1}{n} \sum_{i=1}^{n} X_i, \qquad \hat{\sigma}_0^2 = \frac{1}{n} \sum_{i=1}^{n} (X_i - \overline{X}_n)^2$$

- They are the MLEs for μ and σ^2 in this setting.

Theorem

Let $X_1, \dots, X_n$ be a random sample from $\mathcal{N}(\mu, \sigma^2)$. Then $\bar{X}_n$ and $\hat{\sigma}_0^2$ are independent random variables and $\bar{X}_n \sim \mathcal{N}(\mu, \frac{\sigma^2}{n})$, $\frac{nS_n}{\sigma^2} \sim \chi^2_{n-1}$.

Example

Rain from Seeded Clouds

- Simpson, Olsen, and Eden (1975).
- 26 clouds were seeded with silver nitrate to see if they produced more rain than unseeded clouds.
- Unseeded clouds produce mean rainfall of 4 (log scale).
- We are interested in how far the average log-rainfall of the seeded clouds $\hat{\mu}$ is from 4.

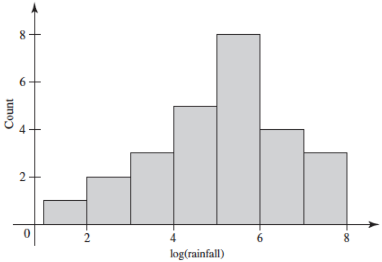

Example

How probable is it that we have overestimated the variance by more than 25%?

$$P(\hat{\sigma}^2 \leq 0.75\sigma^2) = P(\frac{26\hat{\sigma}^2}{\sigma^2} \leq 0.75 * 26) = 0.227$$

What is the smallest number of samples such that

$$\mathrm{P}\left(|\hat{\mu} - \mu| \leq \frac{1}{5}\sigma, \quad |\hat{\sigma} - \sigma| \leq \frac{1}{5}\sigma\right) \geq \frac{1}{2}$$

Example

- $\overline{X}_n = 5.134$, $\hat{\sigma}_0^2 = 63.96/26 = 2.46$
- Let's say I want to answer $P(|\overline{X}_n - \mu| < 5)$.
- If we know σ^2, use CLT.

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

- If we don't know σ^2?

The t distributions

Let $Y \sim \chi^2_m$ and $Z \sim \mathcal{N}(0,1)$ be independent. Then the distribution of $X = \frac{Z}{\left(\frac{Y}{m}\right)^{1/2}}$ is called the t distribution with m degrees of freedom, or t_m.

- PDF of the t distribution:

$$\frac{\Gamma(\frac{m+1}{2})}{(m\pi)^{1/2}\Gamma(\frac{m}{2})}(1+\frac{x^2}{m})^{-(m+1)/2}, -\infty < x < \infty$$

- No closed form CDF, tabulated at the end of statistics books

Relation to the normal distribution

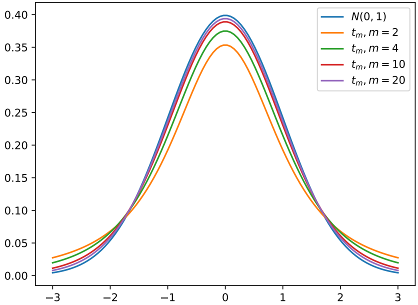

- If $X \sim t_m$ then
  - Moments of the t Distributions:
    - If $m \leq 1$, $E(X)$ does not exist.
    - If $m > 1$, $E(X) = 0$.
    - If $m > 1$, $E\left(|X|^k\right) < \infty$ for $k < m$, $E\left(|X|^k\right) = \infty$ for $k \geq m$.
    - If $(m > 2)$, then $\mathrm{Var}(X) = m/(m-2)$.
  - As $n \to \infty$, t_n converges in pdf to $\mathcal{N}(0,1)$.

Relation to samples of a normal distribution

Theorem (8.4.2)

Let $X_1, \ldots, X_n$ be a random sample from $\mathcal{N}(\mu, \sigma^2)$ and let $\overline{X}_n$ be the sample mean, and define

$$\sigma' = \left(\frac{\sum_{i=1}^{n} (X_i - \overline{X}_n)^2}{n-1} \right)^{1/2}$$

Then $n^{1/2}(\overline{X}_n - \mu)/\sigma')$ follows the t distribution with $n-1$ degrees of freedom.

- Notice that σ' is not the MLE for σ, but $\left(\frac{n-1}{n}\right)^{1/2} \hat{\sigma}_0$
- For large n, $\hat{\sigma}_0$ and σ' are close.

Review

- Let $X_1, \ldots, X_n$ be a random sample from $\mathcal{N}(\mu, \sigma^2)$
- If you know μ but not σ^2

$$\frac{n\hat{\sigma}^2_{MLE}}{\sigma^2} \sim \chi^2_n, \text{ where } \hat{\sigma}^2_{MLE} \text{ is the MLE for } \sigma^2$$

- If you do not know μ or σ^2, then

$$\frac{n\hat{\sigma}^2_0}{\sigma^2} \sim \chi^2_{n-1}, \text{ where } \hat{\sigma}^2_0 = \frac{\sum(X_i - \overline{X}_n)^2}{n} \text{ is the MLE for } \sigma^2$$

$$n^{1/2}(\overline{X}_n - \mu)/\sigma' \sim t_{n-1}, \text{ where } \sigma' = \left(\frac{\sum(X_i - \overline{X}_n)^2}{n-1}\right)^{1/2}$$

Back to our Example

- $\overline{X}_n = 5.134$, $\hat{\sigma}' = \sqrt{63.96/25} = 1.600$
- How confident am I in my $\hat{\mu}$ estimate?
- I know that
$$U = \frac{n^{1/2}(\overline{X}_n - \mu)}{\sigma'} \sim t_{n-1}$$
- I can compute $P(-c < U < c)$.

Confidence Intervals

- I can compute

$$P(\overline{X}_n - \frac{c\sigma'}{n^{1/2}} < \mu < \overline{X}_n + \frac{c\sigma'}{n^{1/2}})$$

Definition (Confidence Interval)

Let $X_1, \ldots, X_n$ be a random sample from $f(x|\theta)$, where θ is unknown. Let $g(\theta)$ be a real-valued function, and let A and B be statistics where $P(A < g(\theta) < B) \geq \gamma \quad \forall\theta$. Then the random interval (A, B) is called a $100\gamma\%$ confidence interval for $g(\theta)$. If equality holds, the CI is exact.

- Notice: A, B are random variables.
- After a random sample is observed, A, B take specific values a and b. The interval (a, b) is then called the observed value of the confidence interval.

Confidence Intervals: Interpretation

- After observing our sample, we find that (a, b) is our 95%-CI for μ.
- This does not mean that $P(a < \mu < b) = 0.95$. In fact, we can not make such statements if we consider μ to be a number (frequentist view).
- We can think of our interpretation as repeated samples.
  - Take a random sample of size n from $\mathcal{N}(\mu, \sigma^2)$.
  - Compute (a, b).
  - Repeat many times.
  - There is a 95% chance for the random intervals to include the value of μ.

Confidence Intervals - the zipper plot

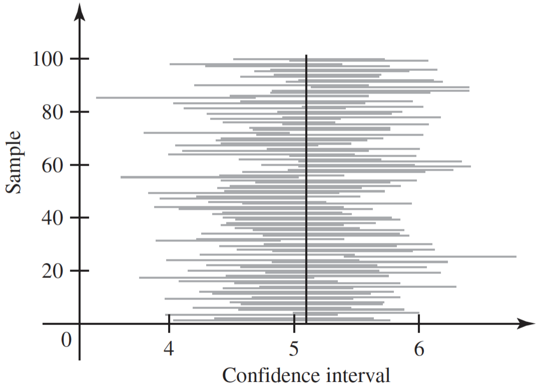

Figure: A sample of one hundred observed 95% confidence intervals based on samples of size 26 from the normal distribution with mean $\mu = 5.1$ and standard deviation $\sigma = 1.6$. In this figure, 94% of the intervals contain the value of μ.

Confidence Intervals

- More generally we want to find $P(c_1 < U < c_2) = \gamma$
- Symmetric confidence intervals: Equal probability on both sides: $P(U \le c_1) = P(U \ge c_2) = \frac{1-\gamma}{2}$
- One-sided confidence interval: All the extra probability is on one side.
- $c_1 = -\infty$ or $c_2 = \infty$.

One-sided Confidence Intervals

Definition (Lower Confidence Limit)

Let A be a statistic so that

$$P(A < g(\theta)) \geq \gamma \quad \forall \theta$$

The random interval (A, ∞) is a one-sided $100\gamma\%$ confidence interval for $g(\theta)$.
A is a $100\gamma\%$ lower confidence limit for $g(\theta)$

Definition (Upper Confidence Limit)

Let B be a statistic so that

$$P(g(\theta) < B) \geq \gamma \quad \forall \theta$$

The random interval $(-\infty, B)$ is a one-sided $100\gamma\%$ confidence interval for $g(\theta)$.
B is a $100\gamma\%$ upper confidence limit for $g(\theta)$

Example

Data on calorie content in 20 different beef hot dogs from Consumer Reports (June 1986 issue):

$$186, 181, 176, 149, 184, 190, 158, 139, 175, 148,$$

$$152, 111, 141, 153, 190, 157, 131, 149, 135, 132$$

- $\overline{X}_n = 156.85$, $\sum_{i=1}^{N}(X_i - \bar{X}_n)^2 = 9740.55$
- Find a 90%-CI for μ.
- Find a lower 90%-CI for μ
- Find a 90%-CI for σ^2.
- If we know that $\sigma^2 = 484$, find a 90%-CI for μ