

From Measurement Uncertainty to Confidence Intervals

Deriving the z and Student- t intervals for a true value

Lecture Notes

Overview

These notes start from the way a measurement result is usually reported,

$$x_{true} = \bar{x} \pm u_x \quad (\text{with probability } \gamma),$$

and turn it into a statement we can actually compute. The route has three stages:

- A. Write the measurement statement as a probability statement and identify what is random.
- B. Assume Gaussian data with known σ and obtain the z -interval.
- C. Replace σ by the sample standard deviation and obtain the Student- t interval.

Each stage uses the same two moves: *standardize* the error so its distribution contains no unknowns, then *read off a quantile*.

1 Stage A: The measurement statement as a probability statement

The statement “ x_{true} lies in $\bar{x} \pm u_x$ with probability γ ” means

$$\bar{x} - u_x \leq x_{true} \leq \bar{x} + u_x,$$

and subtracting $\bar{x}$ and multiplying by -1 rearranges this to

$$P(-u_x \leq \bar{x} - x_{true} \leq u_x) = \gamma. \quad (1)$$

What is random? If $\bar{x}$ were the population mean and x_{true} a fixed constant, the quantity $\bar{x} - x_{true}$ would be a fixed number and (1) would be either 0 or 1 – there is no probability distribution to speak of. The statement becomes meaningful once $\bar{x}$ is a *statistic* computed from data. We therefore replace $\bar{x}$ by the sample mean of N observations,

$$\bar{x}_N = \frac{1}{N} \sum_{i=1}^N x_i,$$

and write μ for the true value (the population mean, assuming no systematic bias):

$$\boxed{P(-u_x \leq \bar{x}_N - \mu \leq u_x) = \gamma} \quad (2)$$

Here μ is a fixed unknown constant and $\bar{x}_N$ is a random variable. The remaining task is to find the u_x that makes (2) true, which requires the distribution of $\bar{x}_N - \mu$.

2 Stage B: Gaussian data, known σ

B1. Assume a model for the data

Let the observations be independent and identically distributed,

$$x_i \stackrel{\text{iid}}{\sim} \mathcal{N}(\mu, \sigma^2), \quad i = 1, \dots, N.$$

B2. Distribution of the sample mean

By linearity of expectation and independence,

$$\mathbb{E}[\bar{x}_N] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}[x_i] = \mu, \quad \text{Var}[\bar{x}_N] = \frac{1}{N^2} \sum_{i=1}^N \text{Var}[x_i] = \frac{\sigma^2}{N}.$$

A linear combination of independent Gaussians is Gaussian, so

$$\bar{x}_N \sim \mathcal{N}\left(\mu, \frac{\sigma^2}{N}\right) \implies \bar{x}_N - \mu \sim \mathcal{N}\left(0, \frac{\sigma^2}{N}\right).$$

The quantity $\sigma / \sqrt{N}$ is the *standard error of the mean*.

B3. Standardize

Divide every term inside (2) by the positive constant $\sigma / \sqrt{N}$; the inequalities are unchanged:

$$P\left(-\frac{u_x}{\sigma/\sqrt{N}} \leq \frac{\bar{x}_N - \mu}{\sigma/\sqrt{N}} \leq \frac{u_x}{\sigma/\sqrt{N}}\right) = \gamma.$$

Define

$$Z = \frac{\bar{x}_N - \mu}{\sigma/\sqrt{N}} \sim \mathcal{N}(0, 1).$$

Z is a *pivot*: a function of the data and of μ whose distribution does not depend on μ or σ .

B4. Read off the quantile

Because Z is standard normal, the value z_γ with $P(-z_\gamma \leq Z \leq z_\gamma) = \gamma$ is known. By symmetry each tail holds $(1 - \gamma)/2$, so

$$z_\gamma = \Phi^{-1}\left(\frac{1 + \gamma}{2}\right),$$

where Φ is the standard normal CDF. Matching the bounds in B3,

$$\frac{u_x}{\sigma/\sqrt{N}} = z_\gamma \implies u_x = z_\gamma \frac{\sigma}{\sqrt{N}}.$$

B5. Unwind

Substituting back and isolating μ :

Result B: known σ

$$\begin{aligned} P\left(-z_\gamma \leq \frac{\bar{x}_N - \mu}{\sigma/\sqrt{N}} \leq z_\gamma\right) &= \gamma \\ \iff P\left(\bar{x}_N - z_\gamma \frac{\sigma}{\sqrt{N}} \leq \mu \leq \bar{x}_N + z_\gamma \frac{\sigma}{\sqrt{N}}\right) &= \gamma \\ u_x = z_\gamma \frac{\sigma}{\sqrt{N}}, \quad z_{95\%} &= 1.960. \end{aligned}$$

3 Stage C: Gaussian data, unknown σ

C1. Estimate σ from the same data

In practice σ is unknown. Use the sample standard deviation

$$s_N^2 = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x}_N)^2.$$

The divisor $N-1$ makes s_N^2 unbiased for σ^2 : one degree of freedom is used up by estimating μ with $\bar{x}_N$.

C2. The naive substitution, and why it is not normal

Replacing σ by s_N in the pivot gives

$$T = \frac{\bar{x}_N - \mu}{s_N/\sqrt{N}}.$$

T is *not* standard normal: the denominator is now random, and when s_N happens to come out small, T is inflated. This produces heavier tails than $\mathcal{N}(0,1)$.

C3. Rewrite T in terms of known distributions

Multiply numerator and denominator by $\sqrt{N}/\sigma$:

$$T = \frac{(\bar{x}_N - \mu)/(\sigma/\sqrt{N})}{s_N/\sigma} = \frac{Z}{\sqrt{V/(N-1)}}, \quad \text{where} \quad V = \frac{(N-1)s_N^2}{\sigma^2}.$$

C4. Two facts about Gaussian samples

For iid Gaussian data (a consequence of Cochran's theorem):

1. $V = \frac{(N-1)s_N^2}{\sigma^2} \sim \chi_{N-1}^2$;
2. $\bar{x}_N$ and s_N^2 are independent, so Z and V are independent.

C5. Identify the distribution

By definition, if $Z \sim \mathcal{N}(0, 1)$ and $V \sim \chi^2_\nu$ are independent, then $Z/\sqrt{V/\nu}$ has a Student- t distribution with ν degrees of freedom. Hence

$$T = \frac{\bar{x}_N - \mu}{s_N / \sqrt{N}} \sim t_{N-1}.$$

The unknown σ has cancelled, so T is again a pivot.

C6. Read off the quantile

Let $t_{\gamma, \nu}$ satisfy $P(-t_{\gamma, \nu} \leq T \leq t_{\gamma, \nu}) = \gamma$, i.e. $t_{\gamma, \nu} = F_{t_\nu}^{-1}((1 + \gamma)/2)$. Then

$$\frac{u_x}{s_N / \sqrt{N}} = t_{\gamma, N-1} \implies u_x = t_{\gamma, N-1} \frac{s_N}{\sqrt{N}}.$$

C7. Unwind

Result C: unknown σ

$$\begin{aligned} P\left(-t_{\gamma, N-1} \leq \frac{\bar{x}_N - \mu}{s_N / \sqrt{N}} \leq t_{\gamma, N-1}\right) &= \gamma \\ \iff P\left(\bar{x}_N - t_{\gamma, N-1} \frac{s_N}{\sqrt{N}} \leq \mu \leq \bar{x}_N + t_{\gamma, N-1} \frac{s_N}{\sqrt{N}}\right) &= \gamma \\ u_x = t_{\gamma, N-1} \frac{s_N}{\sqrt{N}} \end{aligned}$$

This is the original statement (2) with an explicit, computable u_x .

4 The chain at a glance

$$\begin{aligned} P(-u_x \leq \bar{x}_N - \mu \leq u_x) &= \gamma \\ \xrightarrow{x_i \sim \mathcal{N}(\mu, \sigma^2)} \bar{x}_N - \mu &\sim \mathcal{N}(0, \sigma^2/N) \\ \xrightarrow{\div \sigma / \sqrt{N}} P\left(-z_\gamma \leq \frac{\bar{x}_N - \mu}{\sigma / \sqrt{N}} \leq z_\gamma\right) &= \gamma, \quad u_x = z_\gamma \sigma / \sqrt{N} \\ \xrightarrow{\sigma \rightarrow s_N} P\left(-t_{\gamma, N-1} \leq \frac{\bar{x}_N - \mu}{s_N / \sqrt{N}} \leq t_{\gamma, N-1}\right) &= \gamma, \quad u_x = t_{\gamma, N-1} s_N / \sqrt{N} \end{aligned}$$

5 How much does the t correction matter?

The t distribution approaches the standard normal as $\nu \rightarrow \infty$, so $t_{\gamma, N-1} \rightarrow z_\gamma$. Table 1 shows the two-sided 95% quantiles and the factor by which the t interval is wider than the z interval (for the same σ and s_N).

The traditional rule of thumb “use z once $N \gtrsim 30$ ” comes from this table: at $N = 30$ the t interval is only about 4% wider.

N	$\nu = N - 1$	$t_{95\%,\nu}$	$t_{95\%,\nu} / z_{95\%}$
2	1	12.706	6.48
3	2	4.303	2.20
5	4	2.776	1.42
10	9	2.262	1.15
20	19	2.093	1.07
30	29	2.045	1.04
31	30	2.042	1.04
61	60	2.000	1.02
121	120	1.980	1.01
∞	∞	1.960	1.00

Table 1: Two-sided 95% quantiles (0.975 one-sided). $z_{95\%} = 1.960$.




Figure 1: Estimating σ from few samples moves probability into the tails, so a wider interval is needed to capture the same γ .

6 Remarks

Interpretation. In this (frequentist) derivation μ is fixed and the *interval* is random. The statement $P(\cdot) = \gamma$ describes the procedure: over many repetitions of the N -sample experiment, a fraction γ of the computed intervals contain μ . Once a specific interval has been computed, it either contains μ or it does not. A Bayesian treatment with a flat prior on μ gives numerically the same interval but lets you read it as a probability statement about μ itself.

Role of the Gaussian assumption. For the *mean*, the central limit theorem makes $\bar{x}_N$ approximately normal for moderate N even when the x_i are not Gaussian, so Results B and C are fairly robust. The exact χ^2 and independence facts in C4, however, hold only for Gaussian data, which is why the t result is exact only under that assumption.

Scope. The identification $E[\bar{x}_N] = x_{true}$ assumes no systematic error. The interval above quantifies only random (repeatability) scatter – a GUM *Type A* evaluation. Systematic contributions (calibration, instrument offset, resolution, environmental effects) do not shrink with N and are handled as *Type B* components, combined with the *Type A* term to form the total uncertainty.