One-Sample t-Confidence Interval

Estimates a population mean μ. We can be confident that μ is between the upper and lower bound from the formula.

$$\bar{x} \pm t^* \frac{s}{\sqrt{n}}$$

- $\bar{x}$ is the sample mean
- n is the sample size
- s is the sample standard deviation
- t^* is the critical t-value

How to Find the Critical t-Value

In the t-**Distribution Table**, look at the column for your **confidence level** and the row for the **degrees of freedom** df (which is equal to $n - 1$). The number where they line up is t^*. For example: If $n = 5$ and you want a 95% confidence interval, then $t^* = 2.776$.

Confidence Level	60%	80%	90%	95%	98%	99%	99.8%	99.9%
df								
1	1.376	3.078	6.314	12.706	31.821	63.657	318.309	636.619
2	1.061	1.886	2.920	4.303	6.965	9.925	22.327	31.599
3	0.978	1.638	2.353	3.182	4.541	5.841	10.215	12.924
4	0.941	1.533	2.132	2.776	3.747	4.604	7.173	8.610
5	0.920	1.476	2.015	2.571	3.365	4.032	5.893	6.869

If your degrees of freedom aren't on the table, then round down to the nearest df that is.

Assumptions

There are two:

1. **Randomness** You have a simple random sample from a large population (no sample bias).
2. **Normality** The population has a normal distribution.

Even if the population is not normal, t-distribution methods are **robust** when the sample size is large. Here is some guidance on how much to trust these formulas based on the sample size.

- **Small sample** ($n < 30$). Risky. Only use if you are sure the population is roughly normally distributed and the sample has no outliers and very little skew.
- **Large sample** ($n \geq 30$). Probably okay, as long as the data doesn't have any extreme outliers.

If the assumptions are not true, then you should say so. That way people know that your conclusions are not as certain as they might seem.