One Sample T-Tests

Teaching Demo

Ari Smith

Overview

- Hypothesis tests overview
  - Understanding p-values
- One sample T-tests
- Example problem from T-test's origins
  - Interpreting results

Hypothesis Testing: Overview

Hypothesis Test Premise

- Suppose we have n samples of a random variable $X_1, X_2, X_3 \ldots X_n$
- We have a hypothesis H about the model that generates our sample
  - Ex: X drawn from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 1$
- If we calculate a *test statistic* t of our sample, we can evaluate likelihood of a statistic like t being generated if H is true
  - Ex: *sample mean,* $\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$ is a good test statistic if we want to test a hypothesis about distribution mean μ

Hypothesis Test Premise

- We cannot prove H, only reject it or fail to reject it
- Establish a *null hypothesis* H_0, that we initially assume to be true
  - Ex. $\mu = 0$
- Establish an *alternative hypothesis* H_a that contradicts H_0
  - $\mu \neq 0$, $\mu > 0$, or $\mu < 0$ all valid depending on what alternative we're testing (1 tail vs 2 tail)
- Given a *significance level* α, we can compute a *rejection region*, which tells us a range of test statistic values that have a likelihood below α of being observed when H_0 is true
- If our test statistic is in the rejection region, then we can reject H_0 and presume (not prove!) that H_a is true

Steps for Conducting a Hypothesis Test

1. Identify the parameter of interest θ and describe it in the problem context.
2. Determine the null value $\hat{\theta}$ and state the null hypothesis H_0.
3. State the appropriate alternative hypothesis H_a.
4. Give the formula for the computed value of the test statistic t.
5. State the rejection region for the selected significance level α.
6. Compute any necessary sample quantities, substitute into the formula for the test statistic value, and compute that value.
7. Decide if H_0 should be rejected and state this conclusion in the problem context.

p-values from Hypothesis Tests

- Instead of pre-deciding α and calculating rejection region, a *p-value* can be calculated which directly indicates the strength of our evidence against the null hypothesis H_0

- A p-value indicates the likelihood that a test statistic more contradictory to H_0 as t is observed, when H_0 is true

- If $p < \alpha$, then our test statistic is in the rejection region for significance level α

Student’s T-Test

One Sample T-test

- Hypothesis test for the mean of a distribution

- Used in the case of small sample sizes
  - Z-test more common for $n > 30$

Historical Context

- Developed by William Sealy Gosset (1876-1937)
- Head Experimental Brewer at Guinness
- Published anonymously under the name “Student” due to trade secret restrictions
- Interested in hypothesis testing for small sample sizes
  - Farming barley, brewing stout, took lots of time to get samples

Assumptions/Conditions for a T-test

- Assume X is drawn from a normal distribution
- Assume samples drawn from X are independent
- No required knowledge of distribution variance, uses sample variance in calculations instead

"Student's" T distribution

- Under stated assumptions, sample mean test statistic follows a *T distribution* when null hypothesis of $\mu = 0$ is true
- Shape of T distribution depends on *degrees of freedom* ν
  - T distribution approaches normal distribution as ν approaches infinity
- For one sample t-test with n observations, utilize T distribution with $n - 1$ degrees of freedom to calculate rejection regions/p-values

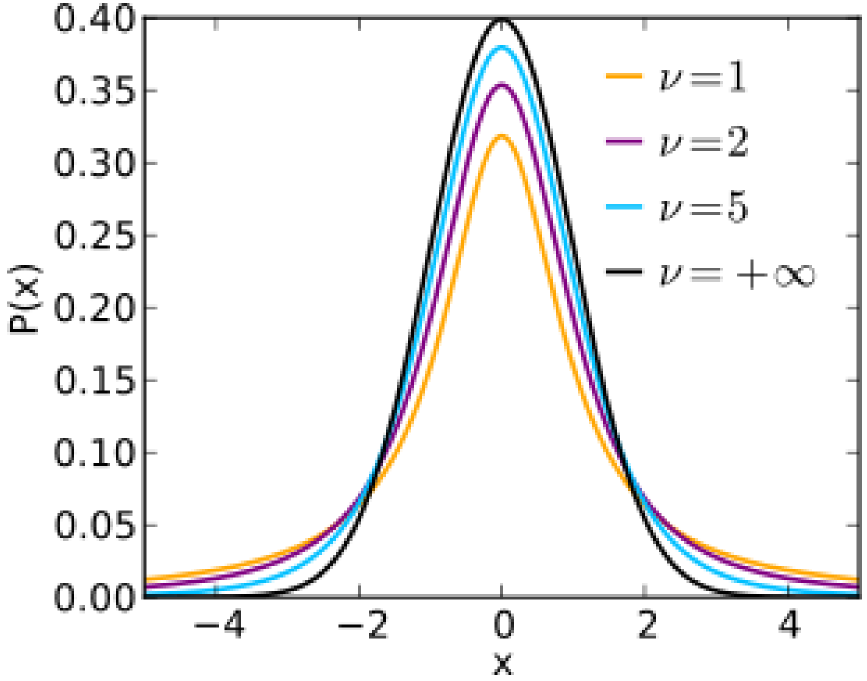

Conducting a One Sample T-test

- Given a sample mean $\bar{X}$, sample standard deviation s, and a hypothesis distribution mean μ_0, the test statistic directly compared against the T distribution is:

$$t = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$$

- Calculating rejection regions from T distributions by hand is *not easy*, lookup tables will give rejection regions for common values of $\alpha,\ \nu$

- Existing packages that can do this for you
  - Excel: TDIST(t, ν, # of tails) — directly calculates p-value
  - Python: scipy.stats.ttest_1samp(sample, μ_0) — calculates t and p-value

Example: Barley Crops

Scenario

- We work for Guinness, making (non-alcoholic) beer from barley farmed on-site
- Barley produces multiple grains per "head"
- The average grain of barley is 42 mg
- Want to determine if our barley's mean grain mass is below average, to see if we should try planting a different variety for better yields
  - Test to significance level of $\alpha = 0.05$

Scenario

- Collected samples of 20 heads of barley from this year's crop
- Data found in a .csv file at: ari-smith-research.github.io/TeachingDemo

Steps for Conducting a Hypothesis Test

1. Identify the parameter of interest θ and describe it in the problem context.
2. Determine the null value $\hat{\theta}$ and state the null hypothesis H_0.
3. State the appropriate alternative hypothesis H_a.
4. Give the formula for the computed value of the test statistic t.
5. State the rejection region for the selected significance level α.
6. Compute any necessary sample quantities, substitute into the formula for the test statistic value, and compute that value.
7. Decide if H_0 should be rejected and state this conclusion in the problem context.

Next Class

- Comparing two populations/samples
- Paired 2 sample T tests: a simple variation on 1 sample
- Unpaired 2 sample T tests
  - Example: comparing two candidate crops

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON