Quantitative Studies: Practical Statistics

Today

- Types of variables
- Correlation
- Hypothesis tests
  - t-test
  - ANOVA
  - Chi-squared test of independence

Variables

- Character (nominal): uninterpreted categories
  - Examples: color; drug type; most experimental treatments
- Numeric (continuous): floating-point numbers
  - Examples: task completion times; test scores; numbers of events
- Numeric (ordinal): items on a scale
  - Examples: Likert scale data

Two Kinds of Activities

- Exploratory data analysis: draw visualizations to help you understand the data
- Statistics: compute numbers to draw inferences

Correlation

- Want to assess strength of association between two variables
- Example: relationship between years of programming experience and typing speed
- EDA: draw a plot

$$r_{xy} = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}}$$

- Pearson's correlation, r^2, represents the fraction of variation in y that is explained by x

Demo

- Autocomplete data: typing speed vs. experience

Hypothesis Tests

- Context: comparing two hypotheses
- Assume H_0 (business as usual)
- Interested in H_1 (something surprising)
- Bias toward H_0: need strong evidence to overturn existing assumptions
- Want: compute $P(H_0)$
  - If less than cutoff (usually 0.05), reject H_0
  - Otherwise, fail to reject H_0

Type I error: incorrectly reject H_0
Type 2 error: incorrectly fail to reject H_0

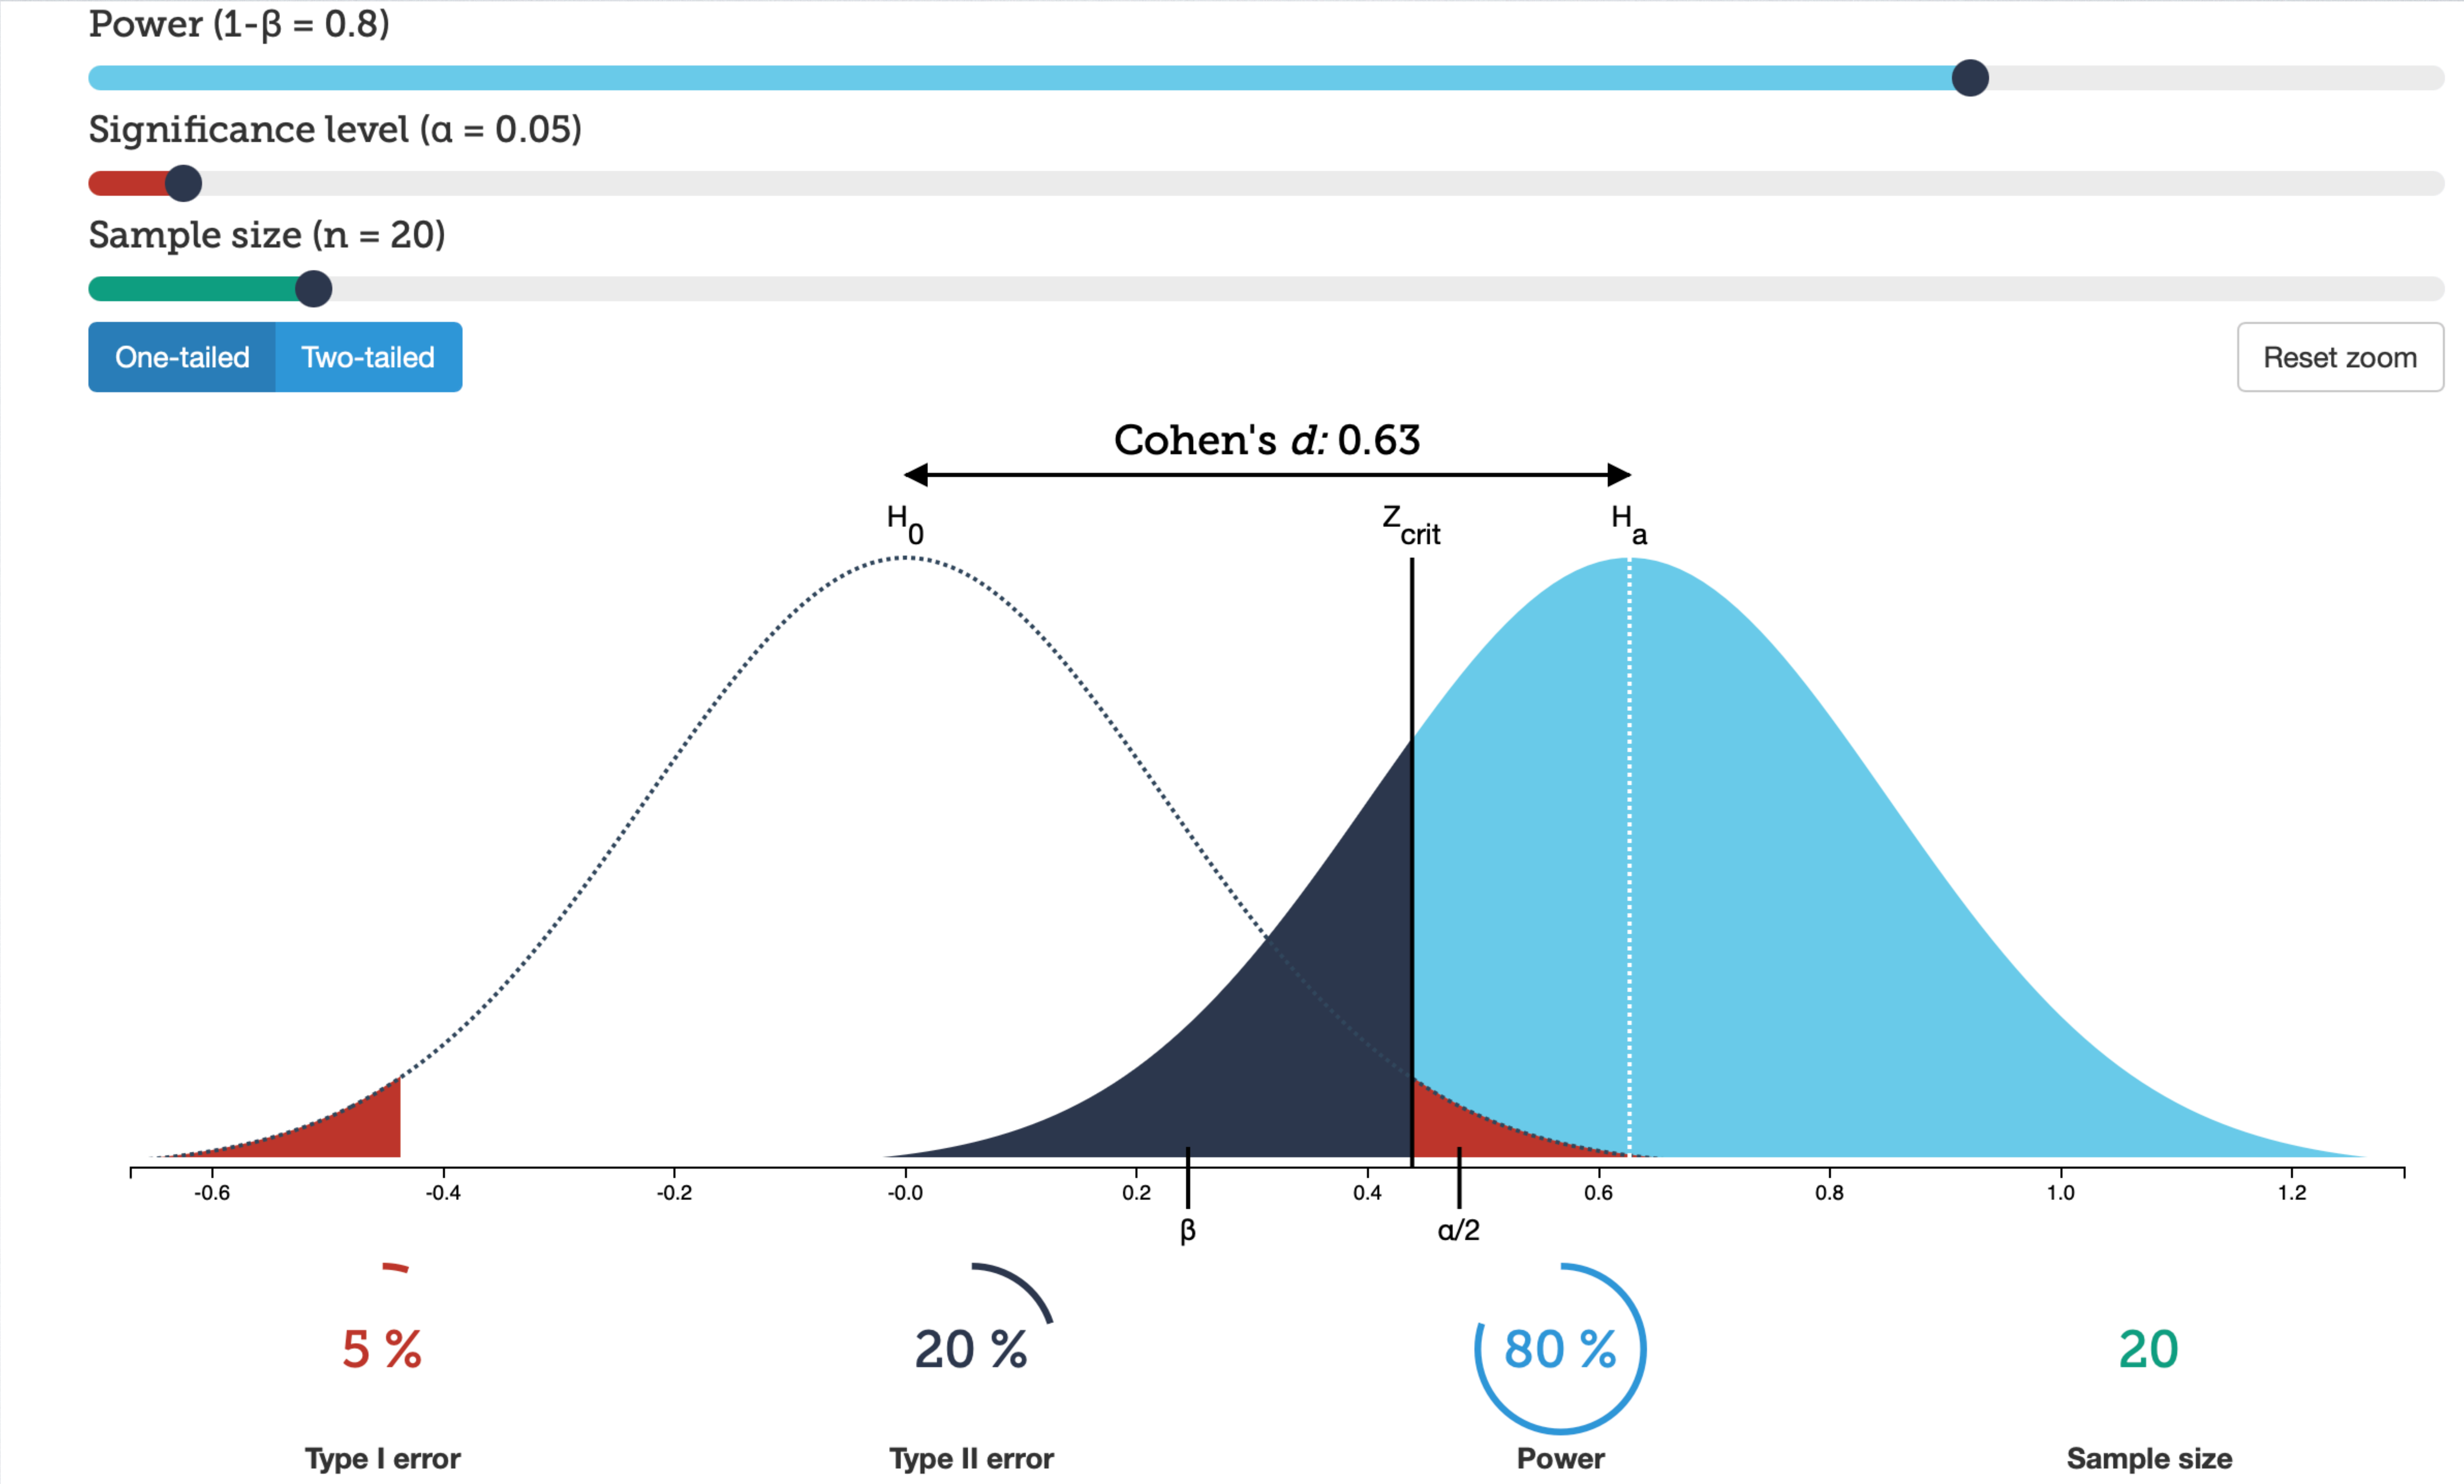

https://rpsychologist.com/d3/nhst/

ANOVA

- Assumptions:
  - variables are **normally distributed**
  - Errors are independent, identically distributed
- Question: do the distributions have the same mean?
- H_0: means are the same
- H_1: means are different
- t test is a special case of ANOVA
- Comes in "repeated measures" (paired) version

ERRORS

- Suppose H_0: $\mu_1 = \mu_2$
- Type 1 error: conclude $\mu_1 \neq \mu_2$ when $\mu_1 = \mu_2$
- Type 2 error: find no significant difference when $\mu_1 \neq \mu_2$
- α: P(type 1 error)
- β: P(type 2 error)

POWER

- Recall: β: P(type 2 error)
- Power: 1 - β
- Probability of rejecting null hypothesis if it is false
- Want more power?
  - Increase N
  - Decrease variance (σ^2)
  - Increase $\mu_1 - \mu_2$

EFFECT SIZE

- Small p-value does not imply a large effect!
- Instead, calculate effect size (Cohen's *d*)
- $d = \frac{\mu_1 - \mu_2}{s}$
- *s*: pooled standard deviation

Interpretation	d
Very small	0.01
Small	0.02
Medium	0.5
Large	0.8
Very large	1.2
Huge	2

Paired Vs. Unpaired

- Paired: for "within" experiments
- Unpaired: for "between" experiments

What if Data Are Not Normally Distributed?

- Then use a nonparametric test
- e.g. Kruskal-Wallis (for unpaired data); Wilcoxon (for paired data)