DSFA

Spring 2019

Lecture 17

Pseudo Random Numbers

Announcements

- Project 2: Posted Monday.
  Due Tuesday, April 9 and April 16

- Prelim 2: In-class. Tuesday, April 16
  (Not Tuesday after spring break)

Cornell University Habitat for Humanity® Campus Chapter

HABITAT FOR HUMANITY BRINGS PEOPLE TOGETHER TO BUILD HOMES, COMMUNITIES, AND HOPE

Interested in Volunteering over Spring Break?

Come join Habitat on our...

SPRING BREAK WORK TRIP

Morgantown, West Virginia

Sign up and find more information here:

habitatcornell.org

Continuous Random Variables

Clicker Question:

Is it possible for someone to be exactly six feet (72 inches) tall?

- A: Yes
- B: No
- C: Impossible to tell

Zehn Deutsche Mark

Uniform Distribution on (a,b)

Informally: All values in the interval (a,b) are equally likely

More formally: Every interval of the same width within (a,b) has the same probability

How to simulate a random sample from a uniform distribution?

Linear Congruential Generator (LCG)

Generate a sequence of integers as follows:

$$X_{n+1} = (aX_n + c) \bmod m$$

where $a = 1103515245\,,\quad c = 12345$

and $m = 2^{31}$. Set $U_{n+1} = X_{n+1}/m$

https://en.wikipedia.org/wiki/Linear_congruential_generator

Does the LCG work?

Does the LCG generate samples that are *indistinguishable* from random samples from the uniform distribution?

Null Hypothesis: The LCG generates random samples from the uniform distribution

Alternative Hypothesis: The LCG does not generate samples from the uniform distribution

Chi-square Statistic

- Divide the unit interval into 10 non-overlapping intervals of equal width 0.1
- Determine the observed and expected counts in each interval
- Calculate test statistic:

$$\chi^2 = \sum \frac{(O - E)^2}{E}$$

Null Distribution!

Cannot use LCG generated values to test the validity of LCG as a uniform random number generator!

Statistical theory shows that the null distribution of the chi-square statistic has a particular form: *a chi-square distribution*

We can compare the LCG-simulated histogram of chi-square values to what the theory predicts

One- versus Two-sided Tests

- **Mendel's Peas**
  Null Hypothesis: Probability of purple flower is 0.75 ($p = 0.75$)
  Alternative Hypothesis: Probability is not 0.75 ($p \neq 0.75$)
  Test Statistic: | p_hat - 0.75 |

  Two-sided test: 'Large' deviation from the null in either direction leads to rejection (of the null hypothesis)

- **Jelly Beans**
  Null Hypothesis: No effect on the probability of acne ($p = 0.2$)
  Alternative Hypothesis: Increase the probability of acne ($p > 0.2$)
  Test Statistic: p_hat - 0.2

  One-sided test: 'Large' positive deviation from the null leads to rejection

Conclusions From a Test

Hypothesis test

Fail to reject the null hypothesis (data is not inconsistent with the null hypothesis - inconclusive)

Reject the null hypothesis (data is inconsistent with the null hypothesis - accept the alternative)

Definition of *P*-value

The P-value is the chance,

- under the null hypothesis,
- that the test statistic
- is equal to the value that was observed in the data or is even further in the direction of the alternative.

Quantifying Conclusions

P(the test statistic would be equal to or more extreme than the observed test statistic under the null hypothesis)

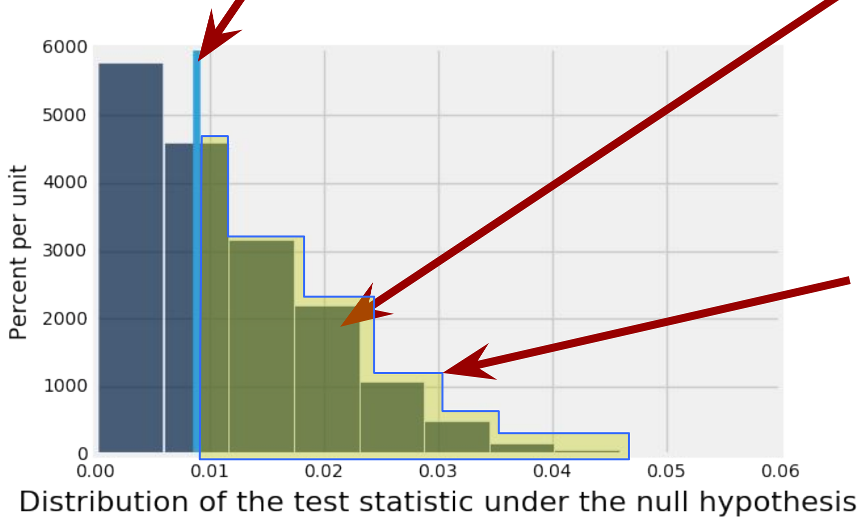

Evaluating Mendel's pea flower hypothesis

This area is the P-value (approximately)

Conventions of Consistency

- **"Inconsistent":** The test statistic is in the tail of the null distribution.
- **"In the tail," first convention:**
  - The area in the tail is less than 5%.
  - The result is "statistically significant."
- **"In the tail," second convention:**
  - The area in the tail is less than 1%.
  - The result is "highly statistically significant."

Can the Conclusion be Wrong?

Yes.

	Null is true	Alternative is true
Test rejects the null	✗	✓
Test doesn't reject the null	✓	✗

(Demo)

An Error Probability

- The cutoff for the P-value is an error probability.
- If:
  - your **cutoff is 5%**
  - and the **null hypothesis happens to be true**
  - (but you don't know that)
- then there is about a **5% chance** that **your test will reject the null hypothesis anyway**.