Lab 4 - The Chi-Square Circus

Overview

Welcome to the Chi-Square Circus, where categories collide, expected counts misbehave, and your AI companion is equal parts genius, trickster, and unreliable narrator.

In this mini lab, you will explore two major applications of the chi-square distribution:

- Testing for **independence**: Are two variables actually connected, or is the universe just messing with you?
- Testing **goodness-of-fit**: Does your data follow a predicted pattern, or does it rebel in unexpected ways?

Your job is to design the scenarios, generate the data, carry out the analyses, and make sense of what unfolds. Your AI's job is to help, complicate, inspire, or derail your thinking, and occasionally, all at once.

AI Use Policy

You may use your AI partner as a brainstorming companion, a generator of unusual ideas, or a source of datasets. However, you are responsible for all reasoning, interpretation, and conclusions. Your work must reflect your own thinking, not a transcript of a conversation.

Deadline: May 27; 4:00pm (Dropbox in Lea)

Activity #1: Independence

Are These Things Actually Related, or Just Awkwardly Appearing Together?.

Guidelines: In this activity, you are testing whether two categorical variables are statistically independent (i.e., whether the patterns in your table could simply be due to chance or whether the variables truly interact within your invented universe). Your task is to imagine a world in which two categorical variables coexist. Your AI assistant may help you generate the setting, categories, characters, or data, but you remain the chief statistician. For example, do enchanted forest creatures prefer certain snacks? Is the choice of study music linked to a student's mood?

Create a two-way table of counts (real or AI-assisted), run the chi-square test for independence, and interpret your results in your chosen universe.

What to Submit

- A short description of your scenario.
- The contingency table you worked with
- Your chi-square test results (statistic, df, and p-value)
- A concise conclusion in context
- One thought-provoking question that invites deeper investigation, curiosity, or mischief about your scenario with solution.
- A short paragraph on how your AI sidekick influenced this investigation.

Activity #2: Goodness of Fit

When Your Expected Pattern Meets Reality (and Reality Does Its Own Thing)

Guideline: In this activity, you are testing whether your observed categorical data actually follow the distribution you claim or expect—whether reality cooperates with your predicted pattern or whether your imagined world stubbornly misbehaves. Your task is to create a categorical variable with at least a few categories and propose an "expected pattern" to test. Your AI assistant may offer ideas or theoretical distributions, but you decide whether those expectations are believable. Examples include a bag of jellybeans that is supposed to contain equal numbers of each colour, or the types of bird calls expected in a calm forest (even though your sample came from a chaotic one), and so on.

Fabricate observed counts. State what distribution you expected. Run a chi-square goodness-of-fit test to evaluate whether your sample plays along.

What to Submit

- A short scenario description

- Observed counts and expected proportions
- Your chi-square test results (statistic, df, and p-value)
- A conclusion in context
- One open-ended question of your own creation that reflects your personal intuition or aesthetic interpretation of your scenario (and solution)