Lab 2 – Probability: Conditions, Dependencies, and the Bayesian Lens

Overview

In this lab, you will move beyond simple counting and explore the rules that govern how events interact. Probability is rarely about isolated incidents; it is about how information changes our degree of certainty. Through a series of interactions with an AI agent, you will investigate the counter-intuitive nature of conditional probability, the ripple effects of the Law of Total Probability, and the power of Bayes' Rule to update our beliefs.

AI Use Policy

As in the previous lab, you are the **Chief Statistician**. The AI is your **simulation engine** and **brainstorming partner**. You are expected to:

- Direct the AI to generate specific datasets or simulations.
- Verify the AI's logic (it is notoriously prone to "hallucinating" probability calculations).
- Perform the formal mathematical proofs and calculations yourself to confirm the AI's empirical findings.

Deadline: April 8th; 4:00 pm

Part 1 – The Monty Hall Simulation

Does information always change the odds?

Task. You will use the AI to simulate the famous "Monty Hall" problem. Your goal is to see if the empirical evidence (the results of many games) matches the theoretical conditional probabilities.

Instructions.

1. Ask the AI to simulate 100 rounds of the Monty Hall game where the contestant **always stays** with their original door. Ask for a summary of the number of wins and losses.
2. Ask the AI to simulate another 100 rounds where the contestant **always switches** to the other remaining door after one is opened.
3. Ask the AI to explain *why* one strategy is better than the other using the language of conditional probability.

What You Need to Submit.

- The results of both simulations (Wins/Losses for "Stay" vs. "Switch").
- A diagram of the sample space and the calculation of $P(\text{Win}|\text{Switch})$.
- A critique of the AI's explanation: Did it correctly identify the role of the host's knowledge? Did it use the Bayes Rule or the role of information? Was the explanation logically sound or did it use circular reasoning? Other weaknesses?
- When first presented with the Monty Hall problem, an overwhelming majority of people assume that each door has an equal probability and conclude that switching does not matter. Pigeons repeatedly exposed to the problem show that they rapidly learn to always switch, unlike humans. Explain what are the majority of humans missing in their assumptions.

Part 2 – The False Positive Paradox

How does a rare event affect the reliability of a test?

Task. In this part, you will construct a scenario involving the **Law of Total Probability** and **Bayes' Rule**. You will investigate why a "highly accurate" test can still be mostly wrong when the event it tests for is rare.

Instructions.

1. Ask the AI to generate a scenario involving a "screening test." This could be a medical test for a rare disease, a security scanner for prohibited items, or a quality control check for a factory defect.

2. The AI must provide:

   - The **Prevalence** of the condition: $P(\text{Condition})$.
   - The **Sensitivity** of the test: $P(\text{Test Positive}|\text{Condition})$.
   - The **Specificity** of the test: $P(\text{Test Negative}|\text{No Condition})$.

3. **Crucial:** Request a scenario where the condition is rare (less than 1% of the population).

What You Need to Submit.

- A description of your scenario and the three probabilities provided by the AI.
- A calculation of the **total probability** of a positive test result, $P(\text{Test Positive})$.
- Using **Bayes' Rule**, calculate the probability that an individual actually has the condition *given* they tested positive: $P(\text{Condition}|\text{Test Positive})$.
- A short written analysis: Compare your calculated result to the "Sensitivity" of the test. Why are they so different? What does this tell us about "Base Rate Neglect"?

Part 3 – Simpson's Paradox and Hidden Variables

Can a treatment appear harmful overall, yet be beneficial for every specific patient group?

Task. In this final section, you will explore the counter-intuitive phenomenon known as **Simpson's Paradox**. You will investigate how marginalizing out a confounding variable can completely reverse the apparent efficacy of a medical treatment, demonstrating why a solid grasp of conditional probability and causal networks is essential in medicine.

Instructions.

1. Ask the AI to generate a hypothetical medical dataset exhibiting Simpson's Paradox. Have it compare two treatments (e.g., Treatment A and Treatment B) for a specific disease.
2. The dataset **must** include a confounding variable, such as the severity of the illness (e.g., Mild vs. Severe cases).
3. Instruct the AI to structure the success/failure numbers such that:
   - Treatment A appears to have a higher *overall* recovery rate than Treatment B.
   - However, when conditioned on the confounder (looking only at Mild cases, and then only at Severe cases), Treatment B has a higher recovery rate in *both* subgroups.

What You Need to Submit.

- The three contingency tables provided by the AI: the aggregated (overall) table, and the two sub-tables separated by illness severity.
- The calculation of the marginal probabilities: $P(\text{Recovery}|\text{Treatment A})$ and $P(\text{Recovery}|\text{Treatment B})$.
- The calculation of the conditional probabilities for both subgroups. For example: $P(\text{Recovery}|\text{Treatment A}, \text{Severe})$ versus $P(\text{Recovery}|\text{Treatment B}, \text{Severe})$.
- A brief written analysis: Use the **Law of Total Probability** to explain the mathematical "trick" causing the paradox. What do the weights (e.g., the proportion of severe cases given Treatment A) reveal about how the treatments were assigned in the first place?
- As an aspiring medical professional or AI researcher, how would mapping a causal graph help prevent you from drawing the wrong clinical conclusion from this purely observational data?