Lecture 17 – Foundations of Probability

DSC 40A, Fall 2022 @ UC San Diego

Mahdi Soleymani, with help from **many others**

Agenda

- Introduction to probability.
- Complement, addition, and multiplication rules.

Introduction to probability

From Lecture 1: course overview

Part 1: Learning from Data (Lectures 1–15)

- Summary statistics and loss functions; mean absolute error and mean squared error.
- Linear regression (incl. linear algebra).
- Clustering.

Part 2: Probability (16–end)

- Set theory and combinatorics; probability fundamentals.
- Conditional probability and independence.
- Naïve Bayes (mix of both parts of the class).

Why do we need probability?

- So far in this class, we have made predictions based on a dataset.
- This dataset can be thought of as a **sample** of some population.
- For a prediction rule to be useful in the future, the sample that was used to create the prediction rule needs to look similar to samples that we'll see in the future.

Probability and statistics

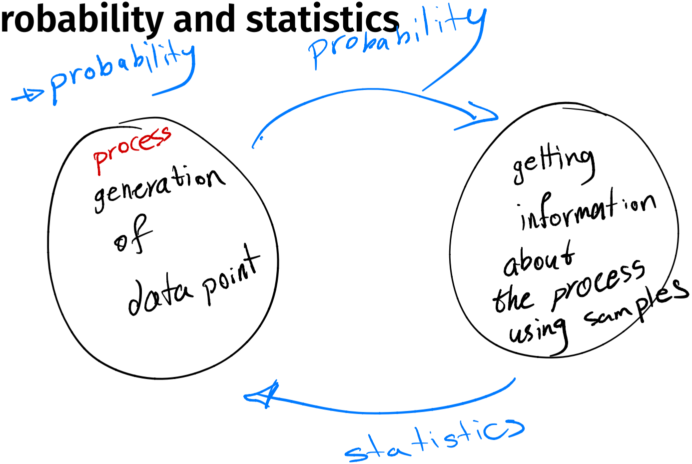

Statistical inference

Given observed data, we want to know how it was generated or where it came from, for the purposes of

- predicting outcomes for other data generated from the same source.
- know how different our sample could have been.
- draw conclusions about our entire population and not just our observed sample (i.e. generalize).

Probability

Given a certain model for data generation, what kind of data do you expect the model to produce? How similar is it to the data you have? Probability is the tool to answer these questions.

- expected value vs. sample mean.
- variance vs. sample variance.
- likelihood of producing exact observed data.

Terminology

yellow 1: 0 0 4 4 4 4 A
y2: 3 3 3 3 3 3 B

- An **experiment** is some process whose outcome is random (e.g. flipping a coin, rolling a die).

Efron's dice

- A **sample space**, S, is the set of all possible outcomes of an experiment.
  - Could be finite or infinite!

$S = \{1, 2, 3, 4, 5, 6\}$

$P(A > B) = 2/3$

$A = \{2, 4, 6\}$

- An **event** is a subset of the sample space.

Example: Rolling a 6-sided die.

$A > B > C > D$

$<$

blue 1 1 1 5 5 5 C
red 2 2 2 2 6 6 D

Probability distributions

outcome: 2, 4

$S = \{1, 2, 3, 4, 5, 6\}$

- A probability distribution, p, describes the **probability** of each outcome s in a sample space S.
  - The probability of each outcome must be between 0 and 1: $0 \le p(s) \le 1$.
  - The sum of the probabilities of each outcome must be exactly 1: $\sum_{s \in S} p(s) = 1$.

$P(1) + P(2) + \cdots + P(6) = 1$

subset of S

- The probability of an **event** is the sum of the probabilities of the outcomes in the event.

$A = \{2, 4, 6\}$

  - $P(E) = \sum_{s \in E} p(s)$. =

$P(A) = P(2) + P(4) + P(6)$

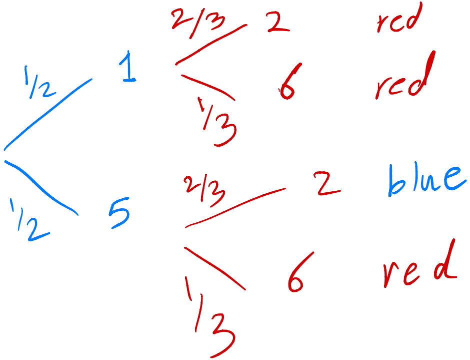

$$\frac{1}{2} \times \frac{2}{3} = \frac{2}{6} = \frac{1}{3}$$

$$\frac{1}{4} \times \frac{1}{3} = \frac{1}{12}$$

$$\frac{1}{4} \times \frac{2}{3} = \frac{1}{6}$$

$$\frac{1}{4} \times \frac{1}{3} = \frac{1}{12}$$

red wins : $\frac{1}{3} + \frac{1}{12} + \frac{1}{12} = \boxed{\frac{2}{3}}$

Complement, addition, and multiplication rules

Sets

- A **set** is an unordered collection of items.
- A **sample space**, S, is the set of all possible outcomes of an experiment.
- An **event**, A is a subset of the sample space. In other words, an event is a set of outcomes.
  - Notationally: $A \subseteq S$.
- $|A|$ denotes the number of elements in set A.

$|A| = \#$ elements in A

Equally-likely outcomes

- If S is a sample space with n possible outcomes, and all outcomes are equally-likely, then the probability of any one outcome occurring is $\frac{1}{n}$.

A: First roll results in H

$A = \{HHH, HHT, HTH, HTT\}$

- The probability of an event A, then, is

$$P(A) = \frac{1}{n} + \frac{1}{n} + \ldots + \frac{1}{n} = \frac{\text{\# of outcomes in A}}{\text{\# of outcomes in S}} = \frac{|A|}{|S|}$$

$P(A) = \frac{4}{8} = \frac{1}{2}$

- **Example:** Flipping a coin three times.

$\frac{1}{|S|} = \frac{1}{8}$

~~$S = \{H, T\}$~~

$S = \{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT\}$

Complement rule

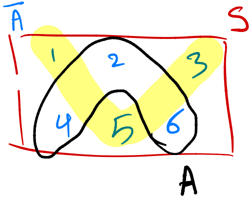

- Let A be an event with probability $P(A)$.
- Then, the event $\bar{A}$ is the **complement** of the event A. It contains the set of all outcomes in the sample space that are not in A.

$A = \{2, 4, 6\} \longrightarrow \bar{A} = \{1, 3, 5\}$

- $P(\bar{A})$ is given by

$$P(\bar{A}) = 1 - P(A)$$

Addition rule

- We say two events are **mutually exclusive** if they have no overlap (i.e. they can't both happen at the same time).

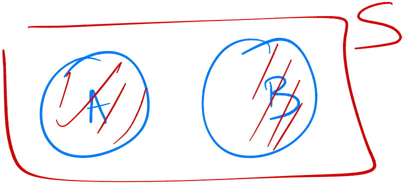

- If A and B are mutually exclusive, then the probability that A or B happens is

"or"

$$P(A \cup B) = P(A) + P(B)$$

$P(A \cup B)$

Principle of inclusion-exclusion

- If events A and B are not mutually exclusive, then the addition rule becomes more complicated.

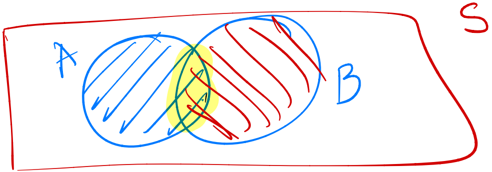

- In general, if A and B are any two events, then

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

A "or" B

Discussion Question

Each day when you get home from school, there is a

- 0.3 chance your mom is at home → A $P(A)=0.3$
- 0.4 chance your brother is at home → B $P(B)=0.4$
- 0.25 chance that both your mom and brother are at home $P(A \cap B)$

When you get home from school today, what is the chance that neither your mom nor your brother are at home?

A) 0.3
B) 0.45
C) 0.55
D) 0.7
E) 0.75

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$
$$= 0.3 + 0.4 - 0.25$$
$$= 0.7 - 0.25 = 0.45$$
$$P(\overline{A \cup B}) = 1 - P(A \cup B) = 0.55$$

To answer, go to `menti.com` **and enter 1372 5551.**

Multiplication rule and independence

- The probability that events A and B both happen is

$$P(A \cap B) = P(A)P(B|A)$$

- $P(B|A)$ is read "the probability that B happens, given that A happened." It is a **conditional probability**.

- If $P(B|A) = P(B)$, events A and B are **independent**.
  - Intuitively, A and B are independent if knowing that A happened gives you no additional information about event B, and vice versa.

  - For two independent events,

$$P(A \cap B) = P(A)P(B)$$

Example: rolling a die

Let's consider rolling a fair 6-sided die. The results of each die roll are independent from one another.

- Suppose we roll the die once. What is the probability that the face is 1 and 2?

- Suppose we roll the die once. What is the probability that the face is 1 or 2?

Example: rolling a die

- Suppose we roll the die 3 times. What is the probability that the face 1 never appears in any of the rolls?

- Suppose we roll the die 3 times. What is the probability that the face 1 appears at least once?

Example: rolling a die

- Suppose we roll the die n times. What is the probability that only the faces 2, 4, and 5 appear?

- Suppose we roll the die twice. What is the probability that the two rolls have different faces?

Summary, next time

Summary

- $\bar{A}$ is the complement of event A. $P(\bar{A}) = 1 - P(A)$.
- Two events A, B are mutually exclusive if they share no outcomes, i.e. they don't overlap. In this case, the probability that A happens or B happens is $P(A \cup B) = P(A) + P(B)$.
- More generally, for any two events, $P(A \cup B) = P(A) + P(B) - P(A \cap B)$.