

Lecture 12

Foundations of Probability

DSC 40A, Spring 2024

Announcements

- Midterm Exam scores are available on Gradescope, and regrade requests are due **tonight**.
- Homework 5 is due on **Thursday at 11:59PM**.
- We're updating the office hours calendar frequently – see dsc40a.com/calendar for the latest.

Agenda

- Overview: Probability and statistics.
- Complement, addition, and multiplication rules.
- Conditional probability.

Note: There are no more DSC 40A-specific readings, but we've posted **many** probability resources on the [resources tab of the course website](#). These will come in handy! Specific resources you should look at:

- The [DSC 40A probability roadmap](#), written by Janine Tiefenbruck.
- The textbook [Theory Meets Data](#), which explains many of the same ideas and contains more practice problems.

Question 🤔

Answer at q.dsc40a.com

Remember, you can always ask questions at q.dsc40a.com!

If the direct link doesn't work, click the "🤔 Lecture Questions"
link in the top right corner of dsc40a.com.

Overview: Probability and statistics

From Lecture 1: Course overview

Part 1: Learning from Data (Weeks 1 through 6)

- Summary statistics and loss functions; empirical risk minimization.
- Linear regression (including multiple variables); linear algebra.
- Clustering.

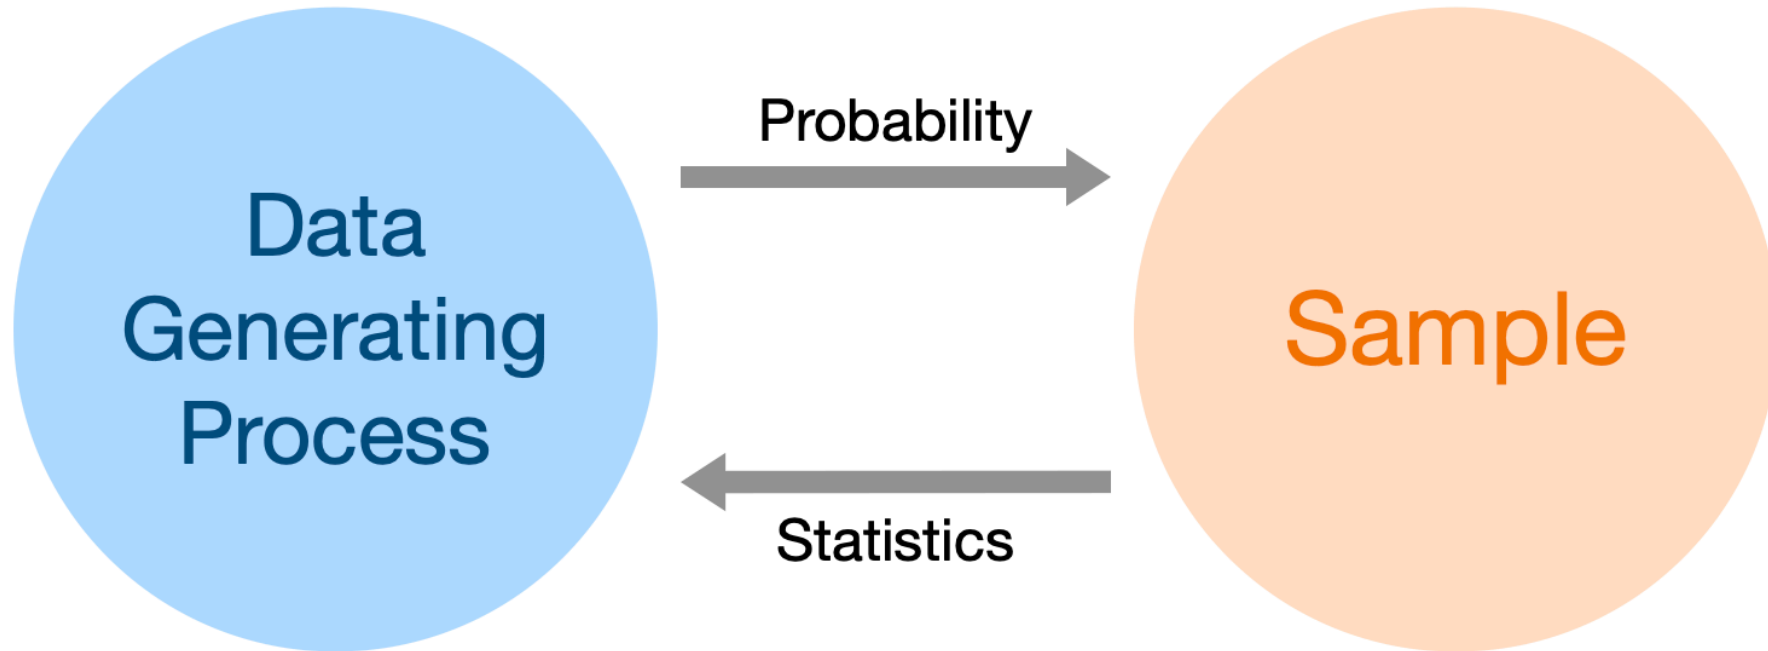
Part 2: Probability (Weeks 7 through 10)

- Set theory and combinatorics; probability fundamentals.
- Conditional probability and independence.
- The Naïve Bayes classifier.

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 hypothesis function to be useful in the future, the sample that was used to create the hypothesis function needs to look similar to samples that we'll see in the future.

Probability and statistics



The plan

- Lecture 12 (today): Key rules of probability.
- Lectures 13-15: Combinatorics.
- Lectures 15-17: Conditional independence and the Naïve Bayes classifier.

Terminology

- An **experiment** is some process whose outcome is random (e.g. flipping a coin, rolling a die).
- A **set** is an unordered collection of items.
 - Sets are usually denoted with { curly brackets }.
 - $|A|$ denotes the number of elements in set A .
- A **sample space**, S , is the set of all possible outcomes of an experiment.
 - Could be finite or infinite!
- An **event** is a subset of the sample space, or a set of outcomes.
 - $E \subseteq S$ means " E is a subset of S ."

Probability distributions

- 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:

$$0 \leq p(s) \leq 1$$

- The sum of the probabilities of each outcome must be exactly 1:

$$\sum_{s \in S} p(s) = 1$$

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

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

What do probabilities *mean*?

- One interpretation: if $\mathbb{P}(E) = p$, then if we repeat our experiment infinitely many times, the proportion of repetitions in which event E occurs is p .
 - If p is large, event E occurs very frequently.

- Another interpretation: $\mathbb{P}(E) = p$ represents our "degree of belief" in the event E .
 - If p is large, we are pretty sure event E is going to happen when we perform our experiment.

Example: Probability of rolling an even number on a six-sided die

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 $\frac{1}{n}$.
- The probability of an event E , then, is:

$$\mathbb{P}(E) = \underbrace{\frac{1}{n} + \frac{1}{n} + \dots + \frac{1}{n}}_{|E| \text{ times}} = \frac{\# \text{ of outcomes in } E}{\# \text{ of outcomes in } S} = \frac{|E|}{|S|}$$

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

Complement, addition, and multiplication rules

Complement rule

- Let A be an event with probability $\mathbb{P}(A)$.
- 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 .

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

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:

$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$$

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:

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

Question 🤔

Answer at q.dsc40a.com

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

- 0.3 chance your mom is at home.
- 0.4 chance your brother is at home.
- 0.25 chance that both your mom and brother are at home.

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

Multiplication rule and independence

- The probability that events A and B both happen is

$$\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B|A)$$

- $\mathbb{P}(B|A)$ means "the probability that B happens, given that A happened." It is a **conditional probability**.
 - More on this soon!
- If $\mathbb{P}(B|A) = \mathbb{P}(B)$, we say 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, $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$.

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 of seeing both 1 and 2?

- Suppose we roll the die once. What is the probability of seeing 1 or 2?

Example: Rolling a die

- Suppose we roll the die 3 times. What is the probability of never seeing a 1 in any of the rolls?

- Suppose we roll the die 3 times. What is the probability of seeing a 1 at least once?

Example: rolling a die

- Suppose we roll the die n times. What is the probability of only seeing the numbers 1, 3, and 4?
- Suppose we roll the die 2 times. What is the probability that the two rolls are different?

Conditional probability

Conditional probability

- The probability of an event may **change** if we have additional information about outcomes.
- Starting with the multiplication rule, $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B|A)$, we have that:

$$\mathbb{P}(B|A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(A)}$$

assuming that $\mathbb{P}(A) > 0$.

Question 🤔

Answer at q.dsc40a.com

Suppose a family has two pets. Assume that it is equally likely that each pet is a dog or a cat. **Consider the following two probabilities:**

- The probability that both pets are dogs given that **the oldest is a dog**.
- The probability that both pets are dogs given that **at least one of them is a dog**.

Are these two probabilities equal?

- A. Yes, they're equal.
- B. No, they're not equal.

Example: Pets

Let's compute the probability that both pets are dogs given that **the oldest is a dog**.


Example: Pets

Let's now compute the probability that both pets are dogs given that **at least one of them is a dog**.

Example: Dominoes

(source: 538)

In a set of dominoes, each tile has two sides with a number of dots on each side: 0, 1, 2, 3, 4, 5, or 6. There are 28 total tiles, with each number of dots appearing alongside each other number (including itself) on a single tile.




Example: Dominoes

Question 1: What is the probability of drawing a "double" from a set of dominoes – that is, a tile with the same number on both sides?


Example: Dominoes

Question 2: Now your friend picks a random tile from the set and tells you that at least one of the sides is a 6. What is the probability that your friend's tile is a double, with 6 on both sides?



Example: Dominoes

Question 3: Now you pick a random tile from the set and uncover only one side, revealing that it has 6 dots. What is the probability that this tile is a double, with 6 on both sides?



[See 538's explanation here.](#)

Summary

- If A is an event, then the complement of A , denoted \bar{A} , is the event that A does not happen, and $P(\bar{A}) = 1 - P(A)$.
- Two events A and B are mutually exclusive if they share no outcomes (no overlap). In this case, $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$.
- More generally, for any two events, $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B)$.
- The probability that events A and B both happen is $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B|A)$.
- $\mathbb{P}(B|A)$ is the **conditional probability** of B occurring, given that A occurs:

$$\mathbb{P}(B|A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(A)}$$

- If $\mathbb{P}(B|A) = \mathbb{P}(B)$, then events A and B are **independent**.