STA 131A: Introduction to Probability Theory

Lecture 2: Set Theory and Probabilistic Models

Dogyoon Song

Spring 2026, UC Davis

Announcements

Office hours:
- Instructor: Wed 2:30–3:30 PM (MSB 4220)
- TA: Tue 3–5 PM (MSB 1117)
- * Please also use Piazza for course questions

Important dates:
- Midterm 1: Fri, Apr 24 (in class)
- Midterm 2: Fri, May 15 (in class)
- Final: Thu, Jun 11, 1:00–3:00 PM
- * *No make-up exams can be arranged other than SDC accommodations*

SDC accommodations: If you need accommodations, please submit requests through the Student Disability Center (SDC) as early as possible

* See the course webpage and syllabus for more details and additional information

Announcements: Homework 1

Homework 1 is now posted

- Due: Tue, April 7, 11:59 PM PT
  - Late submissions will not be accepted for any reason and will receive 0 points

- Submission instructions:
  - You may typeset your solutions in LATEX (preferred), use a word processor, or handwrite and scan them; in all cases, make sure your work is legible and clearly organized
  - Upload a **single PDF file** to Canvas (*Assignments → Homework 1*)
  - Name the file using the prefix of your UC Davis email ID and homework number (e.g., dgsong_hw1.pdf)
  - Please make sure to include "STA 131A," your name, and the last four digits of your student ID on the front page

* See the syllabus for more details about homework policy and requirements

Agenda

- Set theory

- Probabilistic models

- Models vs. reality

Set theory: Notation and terminology

A **set** is a collection of distinct objects, called the **elements** of the set

- $x \in S$: x is an element of S
- $x \notin S$: x is not an element of S

The **empty set** is a set having no elements, denoted by $\emptyset$

Sets can be specified in various ways

- Roster notation (enumeration): $S = \{x_1, x_2, \ldots, x_n\}$, or $S = \{x_1, x_2, \ldots\}$

$$\text{e.g.,} \qquad \{H, T\}, \qquad \{1, 2, 3, 4, 5, 6\}, \qquad \{2, 4, 6, 8, \ldots\}$$

- Set-builder notation (logical formula): $S = \{x \mid x \text{ satisfies } Q\}$

$$\text{e.g.,} \qquad \{2k \mid k \text{ is a positive integer}\}, \qquad \{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$$

Sets may be finite, countably infinite, or uncountable

Set theory: Inclusion relations

Suppose there are two sets S and T

S is a **subset** of T if every element of S is an element of T, i.e., $x \in S \implies x \in T$
- We write $S \subseteq T$ (or $S \subset T$)
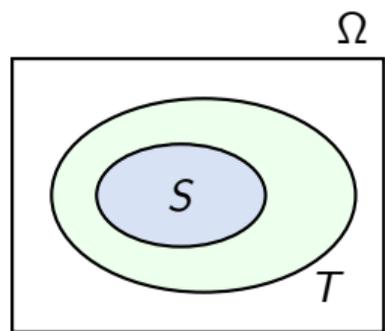- Equivalently, T is a **superset** of S, denoted by $T \supseteq S$ (or $T \supset S$)

The two sets are **equal**, written $S = T$, if $S \subseteq T$ and $T \subseteq S$

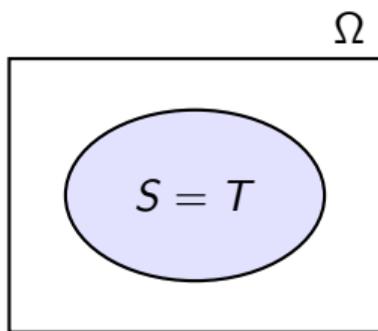S is a **proper (strict)** subset of T, denoted by $S \subsetneq T$, if $S \subseteq T$ and $S \neq T$

* *Remark:* The inclusion relation $\subseteq$ gives a partial order between sets
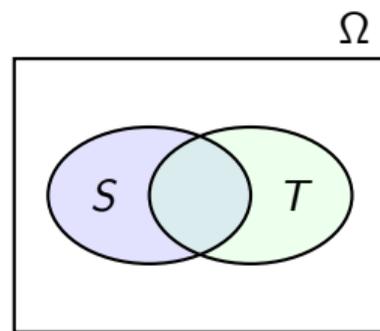
Illustration with Venn diagrams: Inclusion relations

$S \subsetneq T$ $S = T$ Neither set contains the other

Set theory: Set operations

Often, it is convenient to introduce the **universal set** Ω, containing all objects of interest

- $S^c = \{x \in \Omega \mid x \notin S\}$ is the **complement** of S with respect to Ω

Given two sets S and T,

- $S \cup T = \{x \mid x \in S \text{ or } x \in T\}$ is their **union**
- $S \cap T = \{x \mid x \in S \text{ and } x \in T\}$ is their **intersection**

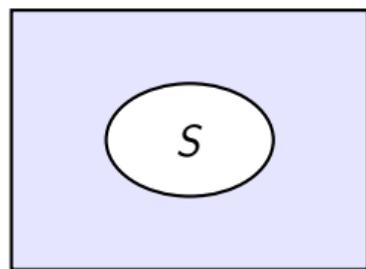Two sets S and T are **disjoint** if $S \cap T = \emptyset$

- A collection of sets is **pairwise disjoint** if no two distinct sets have a common element

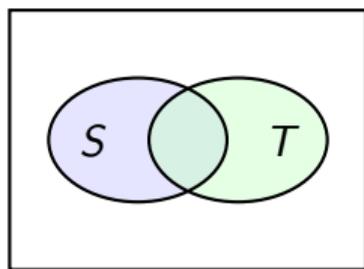A collection of sets $\{S_1, S_2, \ldots\}$ is a **partition** of S if

- $\bigcup_n S_n = S$ and
- the sets $S_1, S_2, \ldots$ are pairwise disjoint

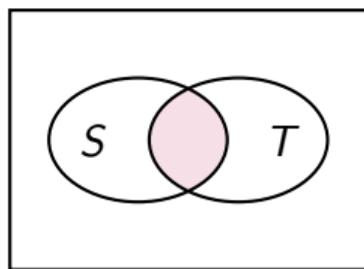# Illustration with Venn diagrams: Set operations
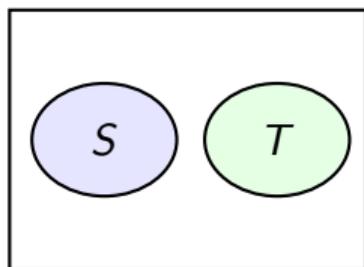

S^c

$S \cup T$

$S \cap T$

Disjoint
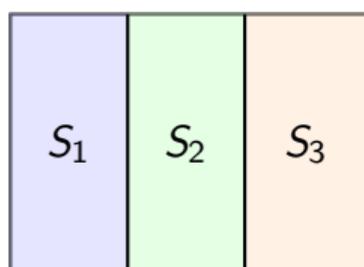
Partition of Ω

Set theory: The algebra of sets

Set operations have several properties, which are direct consequences of the definitions

- Commutativity:
$$S \cup T = T \cup S, \qquad S \cap T = T \cap S$$

- Associativity:
$$S \cup (T \cup U) = (S \cup T) \cup U, \qquad S \cap (T \cap U) = (S \cap T) \cap U$$

- Distributivity:
$$S \cap (T \cup U) = (S \cap T) \cup (S \cap U), \qquad S \cup (T \cap U) = (S \cup T) \cap (S \cup U)$$

- ... and more:
$$(S^c)^c = S, \qquad S \cap S^c = \emptyset, \qquad S \cup S^c = \Omega, \qquad S \cup \Omega = \Omega, \qquad S \cap \Omega = S$$

Set theory: De Morgan's laws

The following properties are known as **De Morgan's laws**:

$$\left(\bigcup_n S_n\right)^c = \bigcap_n S_n^c, \qquad \left(\bigcap_n S_n\right)^c = \bigcup_n S_n^c$$

You can verify these identities, e.g., by using Venn diagrams

However, formally establishing these laws requires a rigorous mathematical proof

- To show $(\bigcup_n S_n)^c \subseteq \bigcap_n S_n^c$:
    - Take $x \in (\bigcup_n S_n)^c$
    - Then $x \notin \bigcup_n S_n$, which implies that for every n, $x \notin S_n$
    - Thus, $x \in S_n^c$ for all n, and therefore, $x \in \bigcap_n S_n^c$
- To show $(\bigcup_n S_n)^c \supseteq \bigcap_n S_n^c$: proceed similarly (*Exercise!*)

Elements of a probabilistic model

A **probabilistic model** is a mathematical description of an uncertain situation

- **Experiment**: an underlying process producing exactly one out of several possible outcomes

A probabilistic models consists of (Ω, P), built to describe an experiment

- **Sample space** Ω: the set of all possible outcomes

    - Event[1]: for this course, think of an event as any subset of Ω

- **Probability law** P: a map that assigns a number $P(A)$ encoding our knowledge or belief about the collective likelihood of the event A, satisfying *certain axioms*

For a legitimate and useful probabilistic model, Ω and P must satisfy some requirements

---

[1]Strictly speaking, some sets have to be excluded, which involves measure theory. However, we can safely ignore pathological issues in this course.

Choosing an appropriate sample space

When choosing a sample space Ω, keep three principles in mind:

- **Collectively exhaustive:** Ω should include every outcome the experiment can produce
    - E.g., Choosing $\Omega = \{1, 2, 3, 4, 5\}$ for a die roll will miss the outcome 6

- **Mutually exclusive:** the outcomes in Ω should be distinguishable from one another
    - E.g., the sample space for a die roll cannot have both "1 or 3" and "1 or 4" as possible outcomes simultaneously

- **No unnecessary detail:** make Ω fine enough to answer the questions of interest, but no finer
    - E.g., including the landing location or maximum height in a coin-toss model is useless

Probability laws

Once the sample space Ω has been chosen, a probability law assigns to each event A a number $P(A)$, called the **probability** of A, subject to three axioms:

1. **Nonnegativity:** $P(A) \geq 0$ for every event A

2. **(Countable) Additivity:** if A and B are disjoint events, then

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

   More generally, if $A_1, A_2, \ldots$ are pairwise disjoint, then

$$P\left(\bigcup_n A_n\right) = \sum_n P(A_n)$$

3. **Normalization:** $P(\Omega) = 1$

Some properties of probability laws

There are many natural properties of a probability law, derivable from the axioms

- $P(\emptyset) = 0$

- $P(A^c) = 1 - P(A)$

- If $A \subseteq B$, then $P(A) \leq P(B)$

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

- $P(A \cup B) \leq P(A) + P(B)$ (union bound)

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

- ...

Constructing a probability law: Example 1

Example (Single coin toss)

Experiment: Tossing a coin once

Possible outcomes: Heads (H) and tails (T)

Sample space: $\Omega = \{H, T\}$

Events: $\{H, T\}, \{H\}, \{T\}, \emptyset$

If the coin is fair—meaning that we know or believe that heads and tails are "equally likely"—then we should assign equal probabilities to the two possible outcomes: $P(\{H\}) = P(\{T\})$

The additivity and the normalization axioms together imply that $P(\{H, T\}) = P(\{H\}) + P(\{T\}) = 1$

Conclusion: The resulting probability law is given by

$$P(\{H, T\}) = 1, \qquad P(\{H\}) = \tfrac{1}{2}, \qquad P(\{T\}) = \tfrac{1}{2}, \qquad P(\emptyset) = 0$$

Constructing a probability law: Example 2

Example (Three coin tosses)

Experiment: Tossing a coin three times

Possible outcomes: length-3 sequences of H's and T's

Sample space: $\Omega = \{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT\}$

Events: subsets of Ω ($2^8 = 256$ subsets)

If the coin is fair, we model the 8 outcomes in Ω as equally likely, so each has probability $1/8$

Consider an event $A = \{\text{exactly 2 heads occur}\}$: the resulting probability law gives

$$
\begin{aligned}
P(A) &= P\big(\{HHT, HTH, THH\}\big) \\
&= P\big(\{HHT\}\big) + P\big(\{HTH\}\big) + P\big(\{THH\}\big) \\
&= \tfrac{1}{8} + \tfrac{1}{8} + \tfrac{1}{8} \\
&= \tfrac{3}{8}
\end{aligned}
$$

Discrete vs. continuous models

In a discrete model, the probabilities of singleton events determine the probability law:

$$P(A) = \sum_{a \in A} P(\{a\}).$$

In a continuous model, singleton probabilities typically do *not* determine the law

Example (A wheel of fortune)

Experiment: Spin a fair wheel and record a number in $[0, 1]$

Possible outcomes: the numbers in the interval $[0, 1]$

Sample space: $\Omega = [0, 1]$

Events: subsets of $[0, 1]$

Q: What is the probability $P(\{x\})$ for a single point $x \in [0, 1]$?

- Under the natural uniform model on $[0, 1]$, $P(\{x\}) = 0$ for every x
- Yet intervals can still have positive probability; for example, $P([a, b]) = b - a$

Models vs. reality

Using the framework of probability theory involves two distinct stages:

1. **Modeling:** Choose a sample space and a probability law to represent reality
   - There are no hard rules for modeling, provided the probability law conforms to the axioms
   - Different reasonable models may represent the same situation differently
   - Even a simplified "incorrect" model may be useful due to simplicity and tractability

2. **Analysis within the model:** Once the model is fixed, compute probabilities and draw conclusions
   - This stage is governed by the axioms of probability and ordinary logic
   - The main difficulty is often computation, not ambiguity

Example: Bertrand's paradox

Apparent paradoxes in probability reflect an ambiguous or incomplete probabilistic model

Example (Bertrand's paradox)

Question: Given a circle and an inscribed equilateral triangle, what is the probability that the length of a randomly chosen chord of the circle is greater than the side of the triangle?
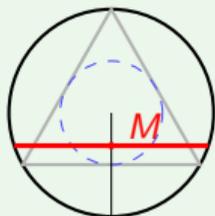
(a) Choose a random point uniformly along radius. A chord is longer than the triangle side iff its midpoint lies within distance $R/2$ of the center. Hence $P = \frac{1}{2}$.
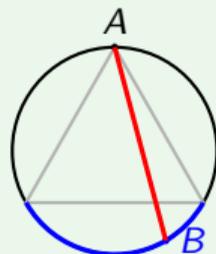
(b) Fix one endpoint A, then choose the other endpoint uniformly on the circumference. The favorable arc has length $1/3$ of the circumference. Hence $P = \frac{1}{3}$.

The phrase *"a randomly chosen chord"* does not specify a unique probabilistic model.

Takeaway: Different models of a "random chord" give different answers.
Probability theory is not inconsistent here.

Wrap-up

Set theory

- Set theory provides a formal language for probability theory; notation and basic operations
- Venn diagrams are useful for intuition, but mathematical proofs come from the definitions

Probabilistic models

- A probabilistic model consists of a sample space Ω and a probability law P
- Choose Ω to be exhaustive, mutually exclusive, and no finer than needed
- P must satisfy the three axioms: nonnegativity, countable additivity, and normalization
- In discrete models, singleton probabilities determine P; in continuous models, they may not

Models vs. reality

- Modeling requires judgment: the same real situation can admit different reasonable models
- Once the model is fixed, probability calculations are governed by the axioms and logic

Suggested reading: [BT08, Ch. 1.1 & 1.2]

References

📄 Dimitri Bertsekas and John N Tsitsiklis.
Introduction to probability, volume 1.
Athena Scientific, 2nd edition, 2008.