COMPSCI 240: Reasoning Under Uncertainty

Andrew Lan and Nic Herndon

University of Massachusetts at Amherst

Spring 2019

Lecture 2: Probability

Recap: Sets

- A set is a collection of objects, which are the elements of the set
- $x \in S$, $x \notin S$, empty set $\emptyset$, number of elements in a set $|S|$
- Subset: $S \subset T$
- Universal set Ω, set complement: $S^c = \{x \in \Omega | x \notin S\}$
- Set union $S \cup T = \{x | x \in S \text{ or } x \in T\}$, intersection $S \cap T = \{x | x \in S \text{ and } x \in T\}$
- **Power Set**: Set of all subsets
- Disjoint sets $S \cap T = \emptyset$
- Partition of a set: S_i and S_j are disjoint for any $i \neq j$ and $S_1 \cup S_2 \cup \cdots \cup S_n = S$

Model of Probability

A probabilistic model is a mathematical description of an uncertain situation. Two fundamental elements of a probabilistic model are

- **Sample Space** Ω: all possible outcomes of an experiment

Model of Probability

A probabilistic model is a mathematical description of an uncertain situation. Two fundamental elements of a probabilistic model are

- **Sample Space** Ω: all possible outcomes of an experiment
- **Probability Law**:

$$A \subset \Omega; \quad P(A),$$

  where A is an **event** (a set of possible outcomes) and $P(A)$ is a non-negative number presenting the **likelihood** of observing the event A.

Probabilistic model involves an **experiment**, which produces an **event** from the **sample space**.

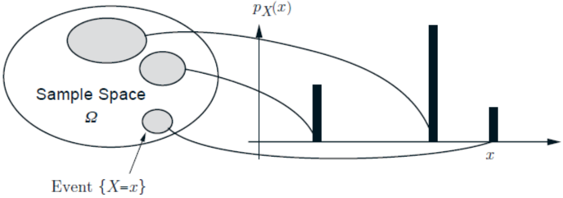

Example

Consider the dice problem. What are the

- **Experiment**:
- **Sample Space**:
- **Possible Outcomes**:

Probability Laws

- Probability represents likelihood of any outcomes or of any set of possible outcomes.
- The probability law assigns to every event A, a number $P(A)$, call the **probability** of A.

Axioms of Probability

- Nonnegativity:

$$P(A) \geq 0$$

- Additivity: For any two disjoint sets A and B,

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

Holds for infinitely many disjoints events $A_1, A_2, A_3, \ldots$

$$P(\cup_i A_i) = \sum_i P(A_i).$$

- Normalization:

$$P(\Omega) = 1$$

That's all we need

Question:

- What is $P(\emptyset)$?
- Can you show that $P(A^c) = 1 - P(A)$?
- If $A \subset B$, then show that $P(A) \leq P(B)$.
- Show that $P(A \cup B) = P(A) + P(B) - P(A \cap B)$
- Sub-additivity: $P(A \cup B) \leq P(A) + P(B)$

Discrete Probability Models

If Ω consists of a finite number of possible outcomes, we are dealing with discrete probability models.
For example,

Discrete Probability Models

If Ω consists of a finite number of possible outcomes, we are dealing with discrete probability models.
For example,

- Coin Toss
- Dice Rolling

Discrete Probability Models

If Ω consists of a finite number of possible outcomes, we are dealing with discrete probability models.
For example,

- Coin Toss
- Dice Rolling

What could be Non-Discrete Probability Models?

Discrete Probability Models

If Ω consists of a finite number of possible outcomes, we are dealing with discrete probability models.
For example,

- Coin Toss
- Dice Rolling

What could be Non-Discrete Probability Models?
For discrete probabilistic models, the probability law is specified by the probabilities of the events that consists of a single element (that are disjoint by nature).

$$A = \{s_1, s_2, \ldots, s_n\} \subset \Omega$$

$$P(\Omega) = P(\{s_1, s_2, ..., s_n\}) = P(s_1) + P(s_2) + ... + P(s_n)$$

Example: Single Coin Toss

Consider an experiment involving a single coin toss.
The sample space (all possible outcomes) is

Example: Single Coin Toss

Consider an experiment involving a single coin toss.
The sample space (all possible outcomes) is

$$\Omega = \{H, T\}$$

All possible events that can occur include

Example: Single Coin Toss

Consider an experiment involving a single coin toss.
The sample space (all possible outcomes) is

$$\Omega = \{H, T\}$$

All possible events that can occur include

$$\{H\}, \{T\}$$

Assuming a fair coin, the associated probability of each event is

Example: Single Coin Toss

Consider an experiment involving a single coin toss.
The sample space (all possible outcomes) is

$$\Omega = \{H, T\}$$

All possible events that can occur include

$$\{H\}, \{T\}$$

Assuming a fair coin, the associated probability of each event is

$$P(\{H\}) = 0.5$$

$$P(\{T\}) = 0.5$$

Verify that the axioms of probability are satisfied!

Recap: Axioms of Probability

- Nonnegativity:

$$P(A) \geq 0 \text{ and } P(B) \geq 0$$

- Additivity: For any two disjoint sets A and B,

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

- Normalization:

$$P(\Omega) = P(A) + P(B) = 1$$

Example: Two Coin Tosses

Now consider an experiment involving two coin tosses.
The sample space (all possible outcomes) is

Example: Two Coin Tosses

Now consider an experiment involving two coin tosses.
The sample space (all possible outcomes) is

$$\Omega = \{HH, HT, TH, TT\}$$

Assuming a fair coin, the associated probability of each outcome is

Example: Two Coin Tosses

Now consider an experiment involving two coin tosses.
The sample space (all possible outcomes) is

$$\Omega = \{HH, HT, TH, TT\}$$

Assuming a fair coin, the associated probability of each outcome is

$$P(\{HH\}) = 0.25$$

$$P(\{HT\}) = 0.25$$

$$P(\{TH\}) = 0.25$$

$$P(\{TT\}) = 0.25$$

Define an event A to observe exactly one head. Then, $P(A)$ is

Example: Two Coin Tosses

Now consider an experiment involving two coin tosses.
The sample space (all possible outcomes) is

$$\Omega = \{HH, HT, TH, TT\}$$

Assuming a fair coin, the associated probability of each outcome is

$$P(\{HH\}) = 0.25$$

$$P(\{HT\}) = 0.25$$

$$P(\{TH\}) = 0.25$$

$$P(\{TT\}) = 0.25$$

Define an event A to observe exactly one head. Then, $P(A)$ is

$$P(A) = 0.5$$

Now, define an event B to observe no head. Then, $P(B)$ is

Example: Two Coin Tosses

Now consider an experiment involving two coin tosses.
The sample space (all possible outcomes) is

$$\Omega = \{HH, HT, TH, TT\}$$

Assuming a fair coin, the associated probability of each outcome is

$$P(\{HH\}) = 0.25$$

$$P(\{HT\}) = 0.25$$

$$P(\{TH\}) = 0.25$$

$$P(\{TT\}) = 0.25$$

Define an event A to observe exactly one head. Then, $P(A)$ is

$$P(A) = 0.5$$

Now, define an event B to observe no head. Then, $P(B)$ is

$$P(B) = 0.25$$

Uniform Discrete Model

If Ω is finite and all possible outcomes are equally likely, it is a **uniform discrete model**. Then, the probability of each element of Ω has the probability of

$$\frac{1}{|\Omega|}$$

More generally, $\forall A \subset \Omega$

$$P(A) = \frac{|A|}{|\Omega|}$$

Uniform Discrete Model - Example

Throwing a fair die is an example of a uniform discrete model.
$\Omega = \{1, 2, 3, 4, 5, 6\}$ Uniform model:

$$P(\{i\}) = \frac{1}{|\Omega|} = \frac{1}{6}$$

for $i = 1, 2, 3, 4, 5, 6$.
A: even number shows up

$$A =$$

Uniform Discrete Model - Example

Throwing a fair die is an example of a uniform discrete model.
$\Omega = \{1, 2, 3, 4, 5, 6\}$ Uniform model:

$$P(\{i\}) = \frac{1}{|\Omega|} = \frac{1}{6}$$

for $i = 1, 2, 3, 4, 5, 6$.
A: even number shows up

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

$$|A| =$$

Uniform Discrete Model - Example

Throwing a fair die is an example of a uniform discrete model.
$\Omega = \{1, 2, 3, 4, 5, 6\}$ Uniform model:

$$P(\{i\}) = \frac{1}{|\Omega|} = \frac{1}{6}$$

for $i = 1, 2, 3, 4, 5, 6$.
A: even number shows up

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

$$|A| = 3$$

$$P(A) =$$

Uniform Discrete Model - Example

Throwing a fair die is an example of a uniform discrete model.
$\Omega = \{1, 2, 3, 4, 5, 6\}$ Uniform model:

$$P(\{i\}) = \frac{1}{|\Omega|} = \frac{1}{6}$$

for $i = 1, 2, 3, 4, 5, 6$.
A: even number shows up

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

$$|A| = 3$$

$$P(A) = \frac{|A|}{|\Omega|} = \frac{3}{6} = \frac{1}{2}.$$

Conditional Probabilities

Conditional probability provides us with a way to reason about the outcome of an experiment based on **partial information** or **observations**.

Conditional Probabilities

Conditional probability provides us with a way to reason about the outcome of an experiment based on **partial information** or **observations**.

Consider rolling a fair die. What is the probability that the outcome is 6 **given that** we know that the outcome is an even number.

- Suppose that you rolled a die while blindfolding yourself. Your friend next to you told you that the number is even. Does that change your probability space?

Conditional Probabilities

Conditional probability provides us with a way to reason about the outcome of an experiment based on **partial information** or **observations**.

Consider rolling a fair die. What is the probability that the outcome is 6 **given that** we know that the outcome is an even number.

- Suppose that you rolled a die while blindfolding yourself. Your friend next to you told you that the number is even. Does that change your probability space?

We can express this conditional probability using $P(A|B)$: conditional probability of A given B, where $P(B) > 0$.
In the above example,

- $A = \{$ The outcome is 6 $\}$
- $B = \{$ The outcome is an even number $\}$

Conditional Probabilities

- A new probability space to be defined.
- The universe (sample space) has been changed to B
- The probability has now to be normalized by $P(B)$

New probability space $P(\cdot|B)$

What is the probability of outcome is 6, given that the outcome is even?

New probability space $P(\cdot|B)$

What is the probability of outcome is 6, given that the outcome is even?

$$P(\text{ the outcome is } 6 | \text{ the outcome is even }) = \frac{1}{3}.$$

Why?

New probability space $P(\cdot|B)$

Definition of conditional probability,

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

New probability space $P(\cdot|B)$

Definition of conditional probability,

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

Another example: What is the probability that the outcome is odd, **given that** the outcome is greater than 1.

New probability space $P(\cdot|B)$

- If A and B are disjoint, i.e., $A \cap B = \emptyset$, then $P(A|B) = 0$. Why?

New probability space $P(\cdot|B)$

- If A and B are disjoint, i.e., $A \cap B = \emptyset$, then $P(A|B) = 0$. Why?
    - In the case of disjoint A and B, $A \cap B = \emptyset$.
    - Which means, $P(A \cap B) = 0$. So $P(A|B) = 0$.
- Can you think of two disjoint events in the experiment of two coin tosses?