Lecture 17 – Foundations of Probability

DSC 40A, Fall 2022 @ UC San Diego

Dr. Truong Son Hy, with help from **many others**

Announcements

- Look at the readings linked on the course website!
- Groupwork Relsease Day: Thursday afternoon
  Groupwork Submission Day: Monday midnight
  Homework Release Day: Friday after lecture
  Homework Submission Day: Friday before lecture
- See `dsc40a.com/calendar` for the Office Hours schedule.
- We will grade the midterm soon.

Agenda

- Solution to the Midterm
- Probability: Complement & addition

Solution to the Midterm

Introduction to Probability

Why do we need probability in Machine Learning?

- Probability is one of the foundations of Machine Learning.
- Learning algorithms will make decisions using probability. Classification models must predict a probability of class membership.
- Algorithms are designed using probability (e.g. Naive Bayes or Gaussian Mixture Models, etc.).
- The data we collect/observe can be understood as samples from some probability distribution.

What is probability?

Informally, a probability distribution $p : X \to \mathbb{R}$ over some domain X is a function such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$.

What is probability?

Informally, a probability distribution $p : X \to \mathbb{R}$ over some domain X is a function such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$.

Example:

- We toss a fair coin, what is the probability of getting head and tail?

What is probability?

Informally, a probability distribution $p : X \to \mathbb{R}$ over some domain X is a function such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$.

Example:

- We toss a fair coin, what is the probability of getting head and tail? $p(\text{head}) = p(\text{tail}) = 0.5$
- We toss a fair dice, what is the probability of getting 1, 2, 3, 4, 5, and 6, respectively?

What is probability?

Informally, a probability distribution $p : X \rightarrow \mathbb{R}$ over some domain X is a function such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$.

Example:

- We toss a fair coin, what is the probability of getting head and tail? $p(\text{head}) = p(\text{tail}) = 0.5$
- We toss a fair dice, what is the probability of getting 1, 2, 3, 4, 5, and 6, respectively? $p(i) = \frac{1}{6}, \forall i \in \{1, .., 6\}$
- If we toss the fair coin and the fair dice infinite number of times, what does the frequencies look like?

What is probability?

Informally, a probability distribution $p : X \to \mathbb{R}$ over some domain X is a function such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$.

Example:

- We toss a fair coin, what is the probability of getting head and tail? $p(\text{head}) = p(\text{tail}) = 0.5$
- We toss a fair dice, what is the probability of getting 1, 2, 3, 4, 5, and 6, respectively? $p(i) = \frac{1}{6}, \forall i \in \{1, .., 6\}$
- If we toss the fair coin and the fair dice infinite number of times, what does the frequencies look like? Tends to the uniform distribution.

Probability: Complement & addition

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 (i.e. cardinality) in set 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 $\frac{1}{n}$.

- The probability of an event A, then, is

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

- **Example:** Flip a (fair) coin three times and we record head/tail each time. How large is the sample space S?

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 A, then, is

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

- **Example:** Flip a (fair) coin three times and we record head/tail each time. How large is the sample space S?
  $|S| = 8$.
  If event A is when head appears exactly once during 3 times, what are $|A|$ and $P(A)$?

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

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

- **Example:** Flip a (fair) coin three times and we record head/tail each time. How large is the sample space S?
  $|S| = 8$.
  If event A is when head appears exactly once during 3 times, what are $|A|$ and $P(A)$? $|A| = 3$ and $P(A) = 3/8$.

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.
- $P(\bar{A})$ is given by
$$P(\bar{A}) = 1 - P(A)$$
- **Example:** Flip a coin 3 times again. What is the probability that there is at least 1 tail?

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.

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

- **Example:** Flip a coin 3 times again. What is the probability that there is at least 1 tail?
Suppose A is the event such that there is no tail. Then, $\bar{A}$ is the event that there is at least 1 tail.

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.
- $P(\bar{A})$ is given by
$$P(\bar{A}) = 1 - P(A)$$
- **Example:** Flip a coin 3 times again. What is the probability that there is at least 1 tail?
Suppose A is the event such that there is no tail. Then, $\bar{A}$ is the event that there is at least 1 tail. We have $|A| = 1$ (i.e. head-head-head) and $P(A) = 1/8$. Therefore, $P(\bar{A}) = 7/8$.

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

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

- **Example:** Flip a fair coin 3 times again. A is the event in which head appears **exactly** once. B is the event in which head appears **exactly** twice. What is the probability that head appears once or twice?

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

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

- **Example:** Flip a fair coin 3 times again. A is the event in which head appears **exactly** once. B is the event in which head appears **exactly** twice. What is the probability that head appears once or twice?
  $|A| = 3$ and $P(A) = 3/8$.

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

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

- **Example:** Flip a fair coin 3 times again. A is the event in which head appears **exactly** once. B is the event in which head appears **exactly** twice. What is the probability that head appears once or twice?
  $|A| = 3$ and $P(A) = 3/8$.
  $|B| = 3$ and $P(B) = 3/8$.

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

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

- **Example:** Flip a fair coin 3 times again. A is the event in which head appears **exactly** once. B is the event in which head appears **exactly** twice. What is the probability that head appears once or twice?
  $|A| = 3$ and $P(A) = 3/8$.
  $|B| = 3$ and $P(B) = 3/8$.
  Therefore, $P(A \cup B) = 3/4$. Is there another way to compute this?

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

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

- **Example:** Flip a fair coin 3 times again. A is the event in which head appears **exactly** once. B is the event in which head appears **exactly** twice. What is the probability that head appears once or twice?
  $|A| = 3$ and $P(A) = 3/8$.
  $|B| = 3$ and $P(B) = 3/8$.
  Therefore, $P(A \cup B) = 3/4$. Is there another way to compute this? Yes, via complement rule.