

Interpreting Probability

Probability can be a subtle.

The first (philosophical) question is “What does probability mean?”

- What does it mean to say that “The probability that the coin landed (will land) heads is $1/2$ ”?

Two standard interpretations:

- Probability is *subjective*: This is a subjective statement describing an individual’s feeling about the coin landing heads
 - This feeling can be quantified in terms of betting behavior
- Probability is an *objective* statement about frequency

Both interpretations lead to the same mathematical notion.

1

Choosing the Sample Space

Example 1: We toss a coin. What’s the sample space?

- Most obvious choice: {heads, tails}
- Should we bother to model the possibility that the coin lands on edge?
- What about the possibility that somebody snatches the coin before it lands?
- What if the coin is biased?

Example 2: We toss a die. What’s the sample space?

Example 3: Two distinguishable dice are tossed together. What’s the sample space?

- $(1,1), (1,2), (1,3), \dots, (6,1), (6,2), \dots, (6,6)$

What if the dice are indistinguishable?

Example 4: You’re a doctor examining a seriously ill patient, trying to determine the probability that he has cancer. What’s the sample space?

Example 5: You’re an insurance company trying to insure a nuclear power plant. What’s the sample space?

3

Formalizing Probability

What do we assign probability to?

Intuitively, we assign them to possible *events* (things that might happen, *outcomes* of an experiment)

Formally, we take a *sample space* to be a *set*.

- Intuitively, the sample space is the set of possible outcomes, or possible ways the world could be.

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

We assign probability to events: that is, to subsets of a sample space.

Sometimes the hardest thing to do in a problem is to decide what the sample space should be.

- There’s often more than one choice
- A good thing to do is to try to choose the sample space so that all outcomes (i.e., elements) are equally likely
 - This is not always possible or reasonable

2

Probability Measures

A *probability measure* assigns a real number between 0 and 1 to every subset of (event in) a sample space.

- Intuitively, the number measures how likely that event is.
- Probability 1 says it’s certain to happen; probability 0 says it’s certain not to happen
- Probability acts like a *weight* or *measure*. The probability of separate things (i.e., disjoint sets) adds up.

Formally, a probability measure \Pr on S is a function mapping subsets of S to real numbers such that:

1. For all $A \subseteq S$, we have $0 \leq \Pr(A) \leq 1$
2. $\Pr(\emptyset) = 0$; $\Pr(S) = 1$
3. If A and B are disjoint subsets of S (i.e., $A \cap B = \emptyset$), then $\Pr(A \cup B) = \Pr(A) + \Pr(B)$.

It follows by induction that if A_1, \dots, A_k are pairwise disjoint, then

$$\Pr(\cup_i^k A_i) = \sum_i^k \Pr(A_i).$$

- This is called *finite additivity*; it’s actually more standard to assume a countable version of this, called *countable additivity*

4

In particular, this means that if $A = \{e_1, \dots, e_k\}$, then

$$\Pr(A) = \sum_{i=1}^k \Pr(e_i).$$

In finite spaces, the probability of a set is determined by the probability of its elements.

Equiprobable Measures

Suppose S has n elements, and we want \Pr to make each element equally likely.

- Then each element gets probability $1/n$
- $\Pr(A) = |A|/n$

In this case, \Pr is called an *equiprobable measure*.

It's not possible in general to put an equiprobable measure on an infinite set.

Theorem: There is no equiprobable measure on the positive integers.

Proof: By contradiction. Suppose \Pr is an equiprobable measure on the positive integers, and $\Pr(1) = \epsilon > 0$.

There must be some N such that $\epsilon > 1/N$. Since $\Pr(1) = \dots = \Pr(N) = \epsilon$, we have

$$\Pr(\{1, \dots, N\}) = N\epsilon > 1 \text{ — a contradiction}$$

But if $\Pr(1) = 0$, then $\Pr(S) = \Pr(1) + \Pr(2) + \dots = 0$.

Examples

Example 1: In the coin example, if you think the coin is fair, and the only outcomes are heads and tails, then we can take $S = \{\text{heads, tails}\}$, and $\Pr(\text{heads}) = \Pr(\text{tails}) = 1/2$.

Example 2: In the two-dice example where the dice are distinguishable, if you think both dice are fair, then we can take $\Pr((i, j)) = 1/36$.

- Should it make a difference if the dice are indistinguishable?

How are the probability of E and \overline{E} related?

- How does the probability that the dice lands either 2 or 4 (i.e., $E = \{2, 4\}$) compare to the probability that the dice lands 1, 3, 5, or 6 ($\overline{E} = \{1, 3, 5, 6\}$)

Theorem: $\Pr(\overline{E}) = 1 - \Pr(E)$.

Proof: E and \overline{E} are disjoint, so that

$$\Pr(E \cup \overline{E}) = \Pr(E) + \Pr(\overline{E}).$$

But $E \cup \overline{E} = S$, so $\Pr(E \cup \overline{E}) = 1$.

Thus $\Pr(E) + \Pr(\overline{E}) = 1$, so

$$\Pr(\overline{E}) = 1 - \Pr(E).$$

Theorem 2: $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$.

$$\begin{aligned} A &= (A - B) \cup (A \cap B) \\ B &= (B - A) \cup (A \cap B) \\ A \cup B &= (A - B) \cup (B - A) \cup (A \cap B) \end{aligned}$$

So

$$\begin{aligned} \Pr(A) &= \Pr(A - B) + \Pr(A \cap B) \\ \Pr(B) &= \Pr(B - A) + \Pr(A \cap B) \\ \Pr(A \cup B) &= \Pr(A - B) + \Pr(B - A) + \Pr(A \cap B) \end{aligned}$$

The result now follows.

Remember the Inclusion-Exclusion Rule?

$$|A \cup B| = |A| + |B| - |A \cap B|$$

This follows easily from Theorem 2, if we take \Pr to be an equiprobable measure. We can also generalize to arbitrary unions.

Conditional Probability

One of the most important features of probability is that there is a natural way to *update* it.

Example: Bob draws a card from a 52-card deck. Initially, Alice considers all cards equally likely, so her probability that the ace of spades was drawn is $1/52$. Her probability that the card drawn was a spade is $1/4$.

Then she sees that the card is black. What should her probability now be that the card is the ace of spades? That it is a spade?

A reasonable approach:

- Start with the original sample space
- Eliminate all outcomes (elements) that you now consider impossible, based on the observation (i.e., assign them probability 0)
- Keep the relative probability of everything else the same.
 - You will need to renormalize to get the probabilities to sum to 1

What should the probability of B be, given that you've observed A ? According to this recipe, it's

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

$$\Pr(A_{\spadesuit}|\text{black}) = (1/52)/(1/2) = 1/26$$

$$\Pr(\text{spade}|\text{black}) = (1/4)/(1/2) = 1/2.$$

A subtlety:

- What if you're not sure that the card is black. How do you take this into account?

Independence

Intuitively, events A and B are independent if they have no effect on each other.

This means that observing A should have no effect on the likelihood we ascribe to B , and similarly, observing B should have no effect on the likelihood we ascribe to A .

Thus, if $\Pr(A) \neq 0$ and $\Pr(B) \neq 0$ and A is independent of B , we would expect

$$\Pr(B|A) = \Pr(B) \text{ and } \Pr(A|B) = \Pr(A).$$

Interestingly, one implies the other.

$$\Pr(B|A) = \Pr(B) \text{ iff } \Pr(A \cap B) / \Pr(A) = \Pr(B) \text{ iff}$$

$$\Pr(A \cap B) = \Pr(A) \times \Pr(B).$$

Formally, we say A and B are (*probablistically*) *independent* if

$$\Pr(A \cap B) = \Pr(A) \times \Pr(B).$$

This definition makes sense even if $\Pr(A) = 0$ or $\Pr(B) = 0$.

The Second-Child Problem

Suppose that any child is equally likely to be male or female, and that the sex of any one child is independent of the sex of the other. You have an acquaintance and you know he has two children, but you don't know their sexes. Consider the following four cases:

1. You visit the acquaintance, and a boy walks into the room. The acquaintance says "That's my older child."
2. You visit the acquaintance, and a boy walks into the room. The acquaintance says "That's one of my children."
3. The acquaintance lives in a culture, where male children are always introduced first, in descending order of age, and then females are introduced. You visit the acquaintance, who says "Let me introduce you to my children." Then he calls "John [a boy], come here!"
4. You go to a parent-teacher meeting. The principal asks everyone who has at least one son to raise their hands. Your acquaintance does so.

In each case, what is the probability that the acquaintance's second child is a boy?

- The problem is to get the right sample space

13

Bayes' Theorem

Bayes Theorem: Let A_1, \dots, A_n be mutually exclusive and exhaustive events in a sample space S .

- That means $A_1 \cup \dots \cup A_n = S$, and the A_i 's are pairwise disjoint: $A_i \cap A_j = \emptyset$ if $i \neq j$.

Let B be any other event in S . Then

$$\Pr(A_i|B) = \frac{\Pr(A_i) \Pr(B|A_i)}{\sum_{j=1}^n \Pr(A_j) \Pr(B|A_j)}$$

Proof:

$$B = B \cap S = B \cap (\cup_{j=1}^n A_j) = \cup_{i=1}^n (B \cap A_j).$$

Therefore, $\Pr(B) = \sum_{j=1}^n \Pr(B \cap A_j)$.

Next, observe that $\Pr(B|A_i) = \Pr(A_i \cap B) / \Pr(A_i)$. Thus,

$$\Pr(A_i \cap B) = \Pr(B|A_i) \Pr(A_i).$$

Therefore,

$$\begin{aligned} \Pr(A_i|B) &= \frac{\Pr(A_i \cap B)}{\Pr(B)} \\ &= \frac{\Pr(B|A_i) \Pr(A_i)}{\sum_{j=1}^n \Pr(B \cap A_j)} \\ &= \frac{\Pr(B|A_i) \Pr(A_i)}{\sum_{j=1}^n \Pr(B|A_j) \Pr(A_j)} \end{aligned}$$

15

Probability Trees

Suppose that the probability of rain tomorrow is .7. If it rains, then the probability that the game will be cancelled is .8; if it doesn't rain, then the probability that it will be cancelled is .1. What is the probability that the game will be played?

The situation can be described by a tree:

Similar trees can be used to describe

- Sequential decisions
- Randomized algorithms

14

Example

In a certain county, 60% of registered voters are Republicans, 30% are Democrats, and 10% are Independents. 40% of Republicans oppose increased military spending, while 65% of the Democrats and 55% of the Independents oppose it. A registered voter writes a letter to the county paper, arguing against increased military spending. What is the probability that this voter is a Democrat?

$S = \{\text{registered voters}\}$

$A_1 = \{\text{registered Republicans}\}$

$A_2 = \{\text{registered Democrats}\}$

$A_3 = \{\text{registered independents}\}$

$B = \{\text{voters who oppose increased military spending}\}$

We want to know $\Pr(A_2|B)$.

We have

$$\begin{aligned} \Pr(A_1) &= .6 & \Pr(A_2) &= .3 & \Pr(A_3) &= .1 \\ \Pr(B|A_1) &= .4 & \Pr(B|A_2) &= .65 & \Pr(B|A_3) &= .55 \end{aligned}$$

16

Using Bayes' Theorem, we have:

AIDS

$$\begin{aligned} \Pr(A_2|B) &= \frac{\Pr(B|A_2) \times \Pr(A_2)}{\Pr(B|A_1) \times \Pr(A_1) + \Pr(B|A_2) \times \Pr(A_2) + \Pr(B|A_3) \times \Pr(A_3)} \\ &= \frac{.65 \times .3}{(.4 \times .6) + (.65 \times .3) + (.55 \times .1)} \\ &= \frac{.195}{.49} \\ &\approx .398 \end{aligned}$$

Suppose we have a test that is 99% effective against AIDS. Suppose we also know that .3% of the population has AIDS. What is the probability that you have AIDS if you test positive?

$S = \{\text{all people}\}$ (in North America??)

$A_1 = \{\text{people with AIDS}\}$

$A_2 = \{\text{people who don't have AIDS}\}$ ($A_2 = \overline{A_1}$)

$B = \{\text{people who test positive}\}$

$$\Pr(A_1) = .003 \quad \Pr(A_2) = .997$$

Since the test is 99% effective:

$$\Pr(B|A_1) = .99 \quad \Pr(B|A_2) = .01$$

Using Bayes' Theorem again:

$$\begin{aligned} \Pr(A_1|B) &= \frac{.99 \times .003}{(.99 \times .003) + (.01 \times .997)} \\ &\approx \frac{.003}{.003 + .01} \\ &\approx .23 \end{aligned}$$

Averaging and Expectation

Suppose you toss a coin that's biased towards heads ($\Pr(\text{heads}) = 2/3$) twice. How many heads do you expect to get?

- In mathematics-speak:
What's the *expected number* of heads?

What about if you toss the coin k times?

What's the average weight of the people in this classroom?

- That's easy: add the weights and divide by the number of people in the class.

But what about if I tell you I'm going to toss a coin to determine which person in the class I'm going to choose; if it lands heads, I'll choose someone at random from the first aisle, and otherwise I'll choose someone at random from the last aisle.

- What's the expected weight?

Averaging makes sense if you use an equiprobable distribution; in general, we need to talk about *expectation*.