Bayes' Rule and LTP

CSE 312 Summer 21
Lecture 6

Announcements

Problem Set 2 and Review Summary 1 have been released.

Please start early and come to office hours!

You can take up to 2 late days on an assignment.

Please list your collaborators in the assignment submission.

Today

Bayes' Rule

Law of Total Probability

More Practice

Wonka Bars

Willy Wonka has placed golden tickets on 0.1% of his Wonka Bars.

You want to get a golden ticket. You could buy a 1000-or-so of the bars until you find one, but that's expensive…you've got a better idea!

You have a test – a very precise scale you've bought.

- If the bar you weigh **does** have a golden ticket, the scale will alert you 99.9% of the time.
- If the bar you weigh does not have a golden ticket, the scale will (falsely) alert you only 1% of the time.

If you pick up a bar and it alerts, what is the probability you have a golden ticket?

Wonka Bars

Willy Wonka has placed golden tickets on 0.1% of his Wonka Bars.

If the bar you weigh **does** have a golden ticket, the scale will alert you 99.9% of the time.

If the bar you weigh does not have a golden ticket, the scale will (falsely) alert you only 1% of the time.

If you pick up a bar and it alerts, what is the probability you have a golden ticket?

Which of these is closest to the right answer?

A. 0.1%
B. 10%
C. 50%
D. 90%
E. 99%
F. 99.9%

Fill out the poll everywhere so Kushal knows how long to explain
Go to pollev.com/cse312su21

Conditioning

Let A be the event you get ALERTED

Let B be the event your bar has a ticket.

What conditional probabilities are each of these?

Willy Wonka has placed golden tickets on 0.1% of his Wonka Bars. $\mathbb{P}(B)$

If the bar you weigh **does** have a golden ticket, the scale will alert you 99.9% of the time. $\mathbb{P}(A|B)$

If the bar you weigh does not have a golden ticket, the scale will (incorrectly) alert you 1% of the time. $\mathbb{P}(A|\bar{B})$

If you pick up a bar and it alerts, what is the probability you have a golden ticket? $\mathbb{P}(B|A)$

Reversing the Conditioning

All of our information conditions on whether B happens or not – does your bar have a golden ticket or not?

But we're interested in the "reverse" conditioning. We know the scale alerted us – we know the test is positive – but do we have a golden ticket?

Bayes' Rule

Bayes' Rule

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

Proof of Bayes' Rule

$\mathbb{P}(A|B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$ by definition of conditional probability

Now, imagining we get $A \cap B$ by conditioning on A, we should get a numerator of $\mathbb{P}(B|A) \cdot \mathbb{P}(A)$

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

As required.

Bayes' Rule

Bayes' Rule

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

What do we know about Wonka Bars?

$$0.999 = \frac{\mathbb{P}(B|A) \cdot \mathbb{P}(A)}{.001}$$

Filling In

What's $\mathbb{P}(A)$?

We'll use a trick called "the law of total probability":

Law of Total Probability

Let $A_1, A_2, \dots, A_k$ be a partition of Ω.

A partition of a set S is a family of subsets $S_1, S_2, \dots, S_k$ such that:

$S_i \cap S_j = \emptyset$ for all i, j and

$S_1 \cup S_2 \cup \cdots \cup S_k = S$.

i.e. every element of Ω is in exactly one of the A_i.

Law of Total Probability

Law of Total Probability

Let $A_1, A_2, \ldots, A_k$ be a **partition** of Ω.

For any event E,

$$\mathbb{P}(E) = \sum_{\text{all } i} \mathbb{P}(E|A_i)\mathbb{P}(A_i)$$

Why?

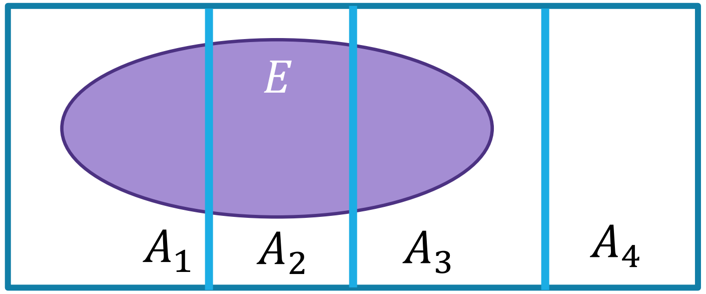

The proof is actually pretty informative on what's going on.

$\sum_{\text{all } i} \mathbb{P}(E|A_i)\mathbb{P}(A_i)$

$= \sum_{\text{all } i} \frac{\mathbb{P}(E \cap A_i)}{\mathbb{P}(A_i)} \cdot \mathbb{P}(A_i)$ (definition of conditional probability)

$= \sum_{\text{all } i} \mathbb{P}(E \cap A_i)$

$= \mathbb{P}(E)$

The A_i partition Ω, so $E \cap A_i$ partition E. Then we just add up those probabilities.

Back to Chocolate

What's $\mathbb{P}(A)$?

We don't know $\mathbb{P}(A)$, but we do know $\mathbb{P}(A|B)$ and $\mathbb{P}\left(A\middle|\bar{B}\right)$. That's a partition of Ω!

$$\mathbb{P}(A) = \mathbb{P}(A|B) \cdot \mathbb{P}(B) + \mathbb{P}\left(A\middle|\bar{B}\right) \cdot P\left(\bar{B}\right)$$

$$= 0.999 \cdot 0.001 + 0.01 \cdot 0.999$$

$$= 0.010989$$

Bayes' Rule

What do we know about Wonka Bars?

$$0.999 = \frac{\mathbb{P}(B|A) \cdot 0.010989}{0.001}$$

Solving $\mathbb{P}(B|A) = \frac{1}{11}$, i.e. about 0.0909.

Only about a 10% chance that the bar has the golden ticket!

Wait a minute…

Willy Wonka has placed golden tickets on 0.1% of his Wonka Bars.
If the bar you weigh **does** have a golden ticket, the scale will alert you 99.9% of the time.
If the bar you weigh does not have a golden ticket, the scale will (falsely) alert you only 1% of the time.

That doesn't fit with many of our guesses. What's going on?

Instead of saying "we tested one and got a positive" imagine we tested 1000. **ABOUT** how many bars of each type are there?

(about) 1 with a golden ticket 999 without. Let's say those are exactly right.

Let's just say that one golden is truly found

(about) 1% of the 999 without would be a positive. Let's say it's exactly 10.

Visually

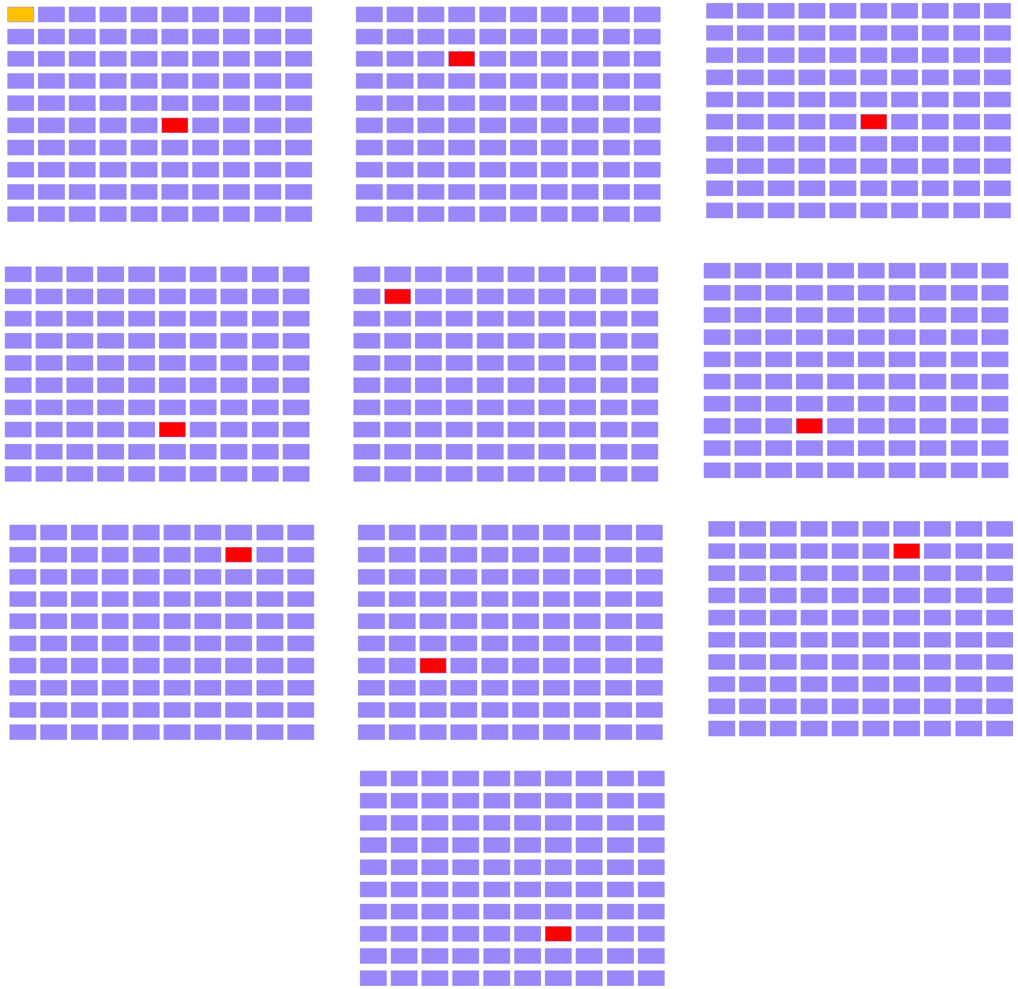

Gold bar is the one (true) golden ticket bar.
Purple bars don't have a ticket and tested negative.
Red bars don't have a ticket, but tested positive.

The test is, in a sense, doing really well.
It's almost always right.

The problem is it's also the case that the correct answer is almost always "no."

Updating Your Intuition

🔥 Take 1: The test is **actually good** and has VASTLY increased our belief that there IS a

If we told you "your job is to find a Wonka Bar with a golden ticket" without the test, you have 1/1000 chance, with the test, you have (about) a 1/11 chance. That's (almost) 100 times better!

This is actually a huge improvement!

Updating Your Intuition

🔥 Take 2: Humans are really bad at intuitively understanding very large or very small numbers.

When I hear "99% chance", "99.9% chance", "99.99% chance" they all go into my brain as "well that's basically guaranteed" And then I forget how many 9's there actually were.

But the number of 9s matters because they end up "cancelling" with the "number of 9's" in the population that's truly negative.

Updating Your Intuition

🔥 Take 3: Viewing tests as updating your beliefs, not revealing the truth.

Bayes' Rule says that $\mathbb{P}(B|A)$ has a factor of $\mathbb{P}(B)$ in it. You have to translate "the test says that there is a golden ticket" to "the test says you should increase your estimate of the chances that you have a golden ticket."

A test takes you from your "prior" beliefs of the probability to your "posterior" beliefs.

More Bayes Practice

Marbles and Coin Tosses

You have three red marbles and one blue marble in your left pocket, and one red marble and two blue marbles in your right pocket.

You will flip a fair coin; if it's heads, you'll draw a marble (uniformly) from your left pocket, if it's tails, you'll draw a marble (uniformly) from your right pocket.

Let B be you draw a blue marble. Let T be the coin is tails.

What is $\mathbb{P}(B|T)$? What is $\mathbb{P}(T|B)$?

Updated Sequential Processes

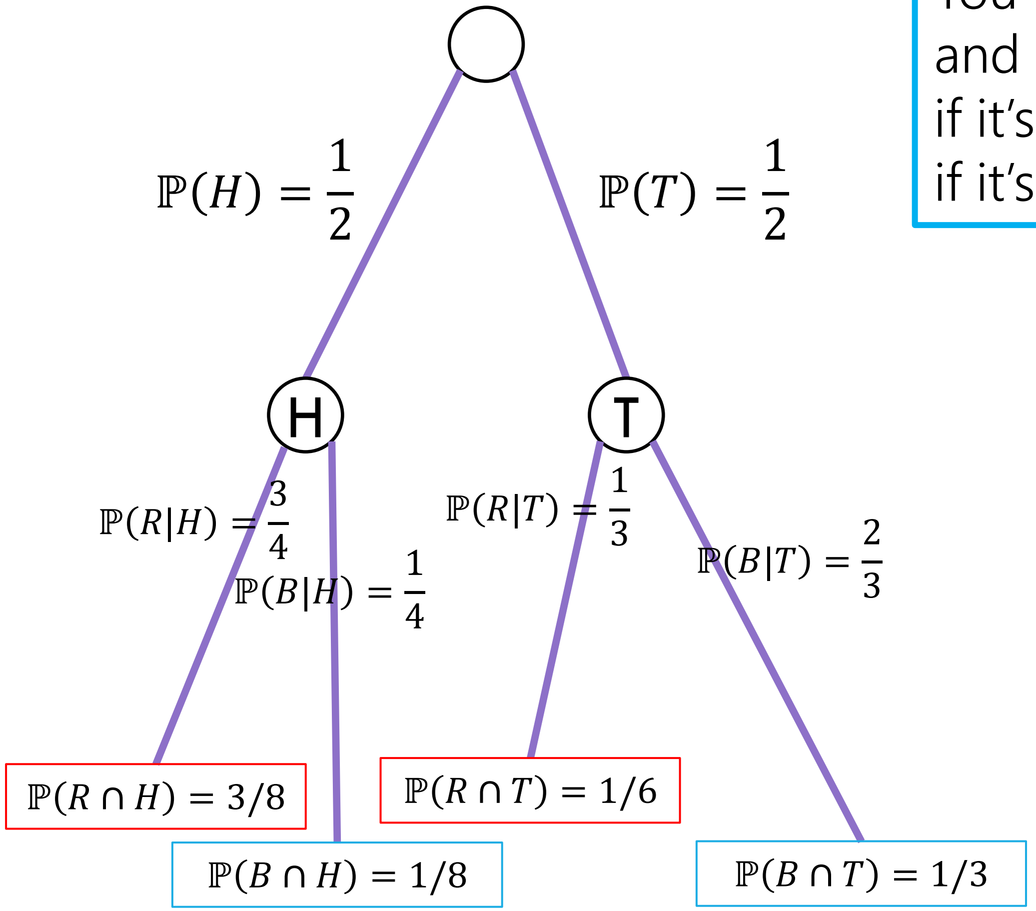

You have three red marbles and one blue marble in your left pocket, and one red marble and two blue marbles in your right pocket.
if it's heads, you'll draw a marble (uniformly) from your left pocket,
if it's tails, you'll draw a marble (uniformly) from your right pocket.

For sequential processes with probability, at each step multiply by
$\mathbb{P}(\text{next step } | \text{all} \cap \text{prior} \cap \text{steps})$

Updated Sequential Processes

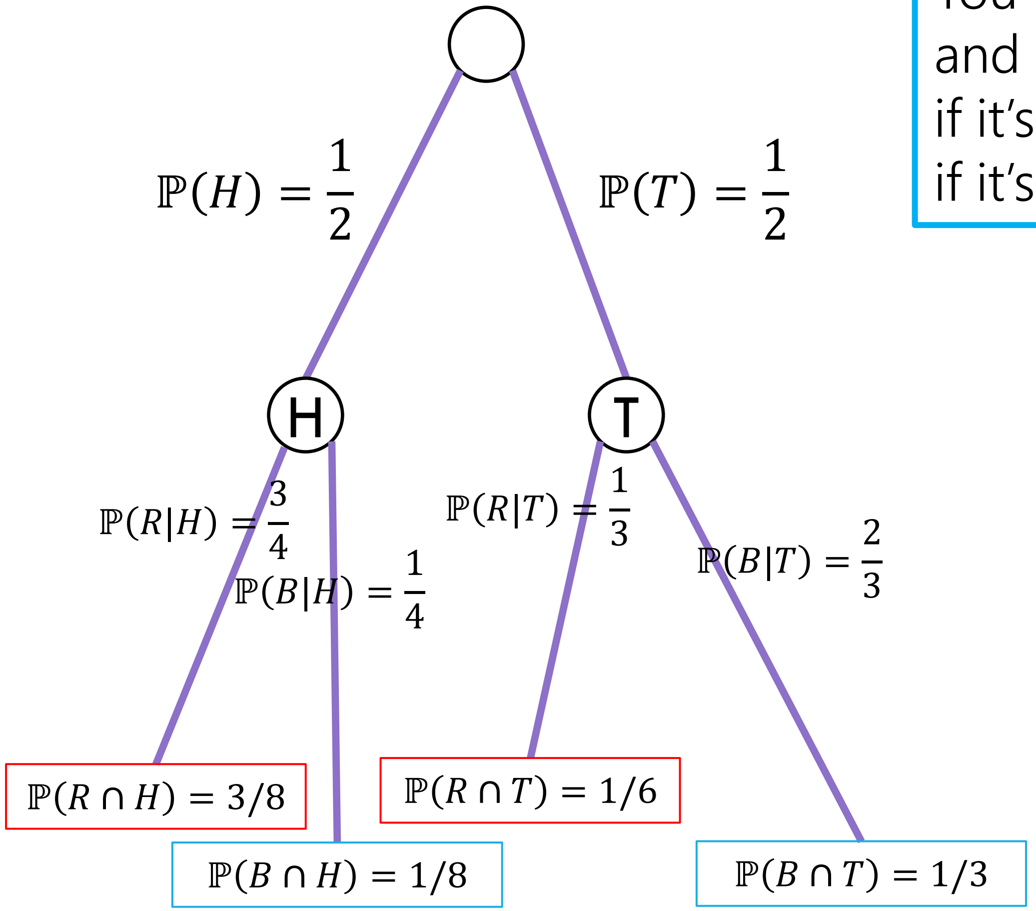

You have three red marbles and one blue marble in your left pocket, and one red marble and two blue marbles in your right pocket.
if it's heads, you'll draw a marble (uniformly) from your left pocket,
if it's tails, you'll draw a marble (uniformly) from your right pocket.

For sequential processes with probability, at each step multiply by
$\mathbb{P}(\text{next step } | \text{all} \cap \text{prior} \cap \text{steps})$

$$\mathbb{P}(B|T) = \frac{2}{3}; \mathbb{P}(B) = \frac{1}{8} + \frac{1}{3} = \frac{11}{24}$$

Flipping the conditioning

What about $\mathbb{P}(T|B)$?

You have three red marbles and one blue marble in your left pocket, and one red marble and two blue marbles in your right pocket.
if it's heads, you'll draw a marble (uniformly) from your left pocket,
if it's tails, you'll draw a marble (uniformly) from your right pocket.

Pause, what's your intuition?

Is this probability

A. less than ½

B. equal to ½

C. greater than ½

Fill out the poll everywhere
pollev.com/cse312su21

Flipping the conditioning

You have three red marbles and one blue marble in your left pocket,
and one red marble and two blue marbles in your right pocket.
if it's heads, you'll draw a marble (uniformly) from your left pocket,
if it's tails, you'll draw a marble (uniformly) from your right pocket.

What about $\mathbb{P}(T|B)$?

Pause, what's your intuition?

Is this probability

A. less than ½

B. equal to ½

C. greater than ½

The right (tails) pocket is far more likely to produce a blue marble if picked than the left (heads) pocket is. Seems like $\mathbb{P}(T|B)$ should be greater than ½.

Flipping the conditioning

What about $\mathbb{P}(T|B)$?

You have three red marbles and one blue marble in your left pocket, and one red marble and two blue marbles in your right pocket.
if it's heads, you'll draw a marble (uniformly) from your left pocket,
if it's tails, you'll draw a marble (uniformly) from your right pocket.

Bayes' Rule says:

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

$$= \frac{\frac{2}{3} \cdot \frac{1}{2}}{\frac{11}{24}} = \frac{8}{11}$$

Some Technical Notes

Technical Note

After you condition on an event, what remains is a probability space.

With B playing the role of the sample space,

$\mathbb{P}(\omega|B)$ playing the role of the probability measure.

All the axioms are satisfied (it's a good exercise to check)

That means any theorem we write down has a version where you condition everything on B.

An Example

Bayes Theorem still works in a probability space where we've already conditioned on S.

$$\mathbb{P}(A|[B \cap S]) = \frac{\mathbb{P}(B|[A \cap S]) \cdot \mathbb{P}(A|S)}{\mathbb{P}(B|S)}$$

A word of caution!

I often see students write things like

$\mathbb{P}([A|B]|C)$

This is not a thing.

You probably want $\mathbb{P}(A|[B \cap C])$

$A|B$ isn't an event – it's describing an event **and** telling you to restrict the sample space. So, you can't ask for the probability of that conditioned on something else.

Extra Practice

Where There's Smoke There's…

There is a dangerous (you-need-to-call-the-fire-department-dangerous) fire in your area 1% of the time.

If there is a dangerous fire, you'll smell smoke 95% of the time;

If there is not a dangerous fire, you'll smell smoke 10% of the time (barbecues are popular in your area)

If you smell smoke, should you call the fire department?

S be the event you smell smoke

F be the event there is a dangerous fire

$$\mathbb{P}(F|S) = \frac{\mathbb{P}(S|F)\cdot\mathbb{P}(F)}{\mathbb{P}(S)} = \frac{\mathbb{P}(S|F)\cdot\mathbb{P}(F)}{\mathbb{P}(S|F)\mathbb{P}(F)+\mathbb{P}\left(S\middle|\bar{F}\right)\mathbb{P}(\bar{F})}$$

$$= \frac{.95\cdot.01}{.95\cdot.01+.1\cdot.99} \approx .088$$

Probably not time yet to call the fire department.