CS70 – Spring 2024

Lecture 16 – March 12

Last Lecture:

- Definition of a probability space:

  Ω = set of outcomes

  $\Pr[\omega]$ = probability for each $\omega \in \Omega$

- Events $E \subseteq \Omega$

$$\Pr[E] = \sum_{\omega \in E} \Pr[\omega]$$

- Uniform probability space:

$$\Pr[\omega] = \frac{1}{|\Omega|} \quad \forall\, \omega \in \Omega$$

$$\Pr[E] = \frac{|E|}{|\Omega|} \quad \forall\, E \subseteq \Omega$$

Ref: Note 13

Today:

- Conditional probability
- Intersections & unions of events
- Bayes Rule & inference

Ref: Note 14

Conditional Probability

Recall : 5-card poker hand

$\rightarrow$ uniform prob. space with $|\Omega| = \binom{52}{5}$

Event E_{Flush} = all five cards of same suit

$$\Pr[E_{Flush}] = \frac{|E_{Flush}|}{|\Omega|} = \frac{4 \times \binom{13}{5}}{\binom{52}{5}} \approx 0.002$$

Now suppose your first 4 cards are all $\diamondsuit$

What is now $\Pr[E_{Flush}]$?

$$\Pr[E_{Flush} \mid \diamondsuit\diamondsuit\diamondsuit\diamondsuit] = \frac{\#\text{remaining } \diamondsuit}{\#\text{remaining cards}} = \frac{9}{48} \approx 0.19$$

Defn: For any events A, B with $\Pr[B] > 0$, the <u>conditional probability of A given B</u> is

$$\Pr[A|B] = \frac{\Pr[A \cap B]}{\Pr[B]}$$

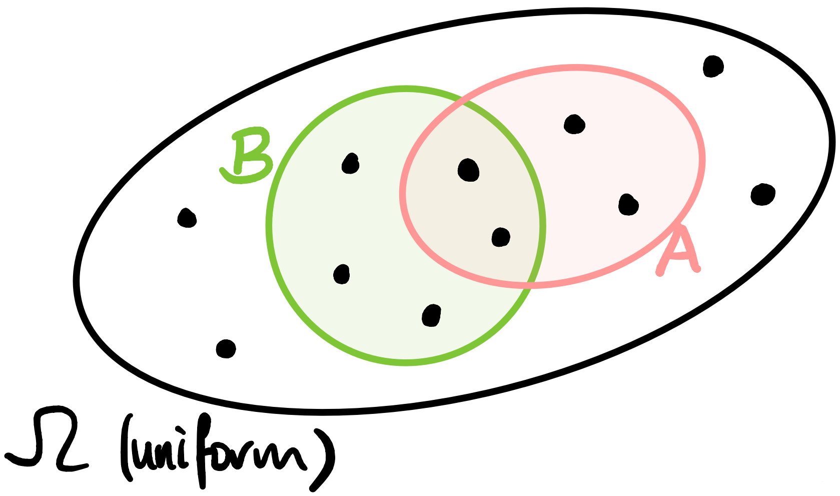

$\Pr[A] = \frac{4}{11}$

$\Pr[A|B] = \boxed{\frac{2}{5}}$

For each sample point $\omega \in B$: $\Pr[\omega] \rightarrow \frac{\Pr[\omega]}{\Pr[B]}$

- - - - - - - - - - $\omega \notin B$: $\Pr[\omega] \rightarrow 0$

Then $\Pr[A] = \sum_{\omega \in A} \Pr[\omega] \longrightarrow \sum_{\omega \in A \cap B} \frac{\Pr[\omega]}{\Pr[B]} = \frac{\Pr[A \cap B]}{\Pr[B]}$

Example: Flush

$|\Omega| = 52!$

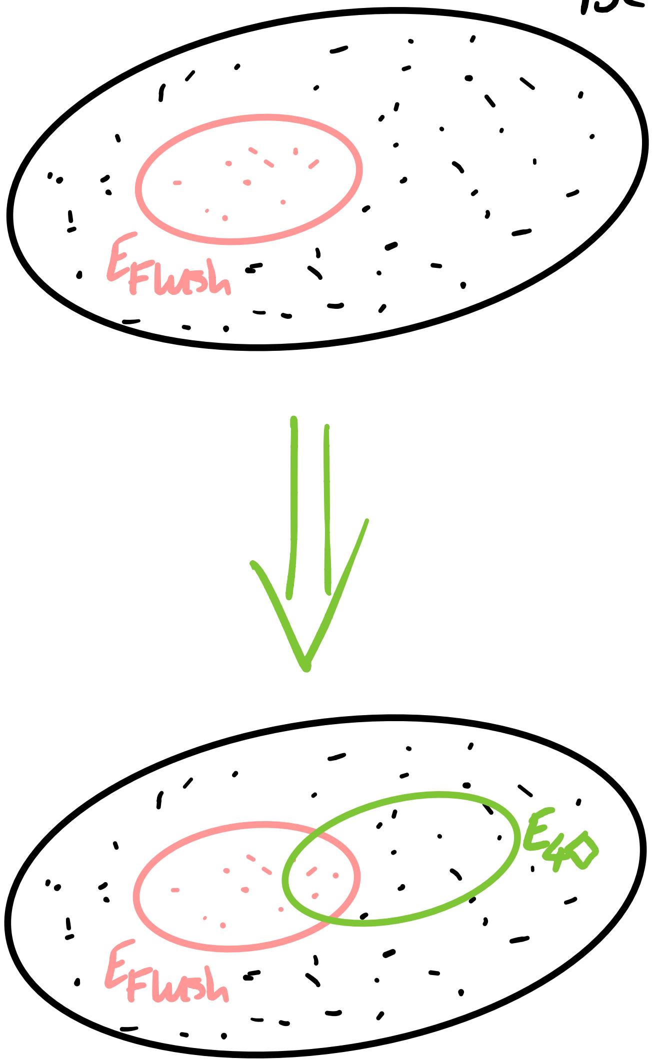

$$\Pr[E_{Flush}] = \frac{|E_{Flush}|}{|\Omega|} \simeq 0.002$$

$$\Pr[E_{Flush} \mid E_{4\diamond}] = \frac{|E_{Flush} \cap E_{4\diamond}|}{|E_{4\diamond}|} = \frac{\binom{13}{5}}{\binom{13}{4} \times 48} = \frac{9}{48}$$

Example: Dice Game

Roll 2 dice – you win if sum is ≥ 9

$$\Pr[\text{win}] = \frac{|W|}{|\Omega|} = \frac{10}{36} = \frac{5}{18}$$

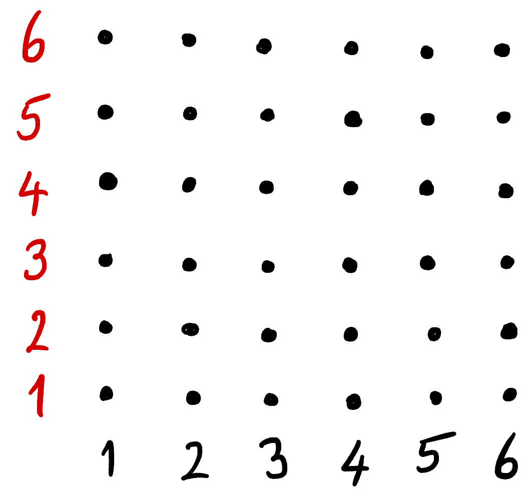

Define E_i = "red die shows i"

$$\Pr[W \mid E_6] = \frac{\Pr[W \cap E_6]}{\Pr[E_6]}$$

$$= \frac{4/36}{1/6} = \boxed{\frac{2}{3}} > \Pr[W]$$

$$\Pr[W \mid E_3] = \frac{\Pr[W \cap E_3]}{\Pr[E_3]}$$

$$= \frac{1/36}{1/6} = \boxed{\frac{1}{6}} < \Pr[W]$$

Example: Coin Tossing

Toss a fair coin 20 times

E_i = "ith toss comes up Heads"

$\Pr[E_i] = 1/2 \quad \forall i$

Suppose the first 19 tosses all come up Heads

What is now $\Pr[E_{20}]$?

$$\Pr[E_{20} \mid E_1 \cap \ldots \cap E_{19}] = \frac{\Pr[E_1 \cap \ldots \cap E_{20}]}{\Pr[E_1 \cap \ldots \cap E_{19}]} = \frac{1/2^{20}}{1/2^{19}}$$

$$= \boxed{\frac{1}{2}} = \Pr[E_{20}]$$

We say that E_{20} is <u>independent</u> of $E_1, \ldots, E_{19}$

Correlation

We have seen that $\Pr[A|B]$ can be $\left\{\begin{matrix} > \\ = \\ < \end{matrix}\right\}$ $\Pr[A]$

$\Pr[A|B] > \Pr[A] \longrightarrow$ A, B <u>positively</u> correlated

$\Pr[A|B] < \Pr[A] \longrightarrow$ A, B <u>negatively</u> correlated

$\Pr[A|B] = \Pr[A] \longrightarrow$ A, B <u>independent</u>

E.g. Uniform prob. space over US population

A = "gets lung cancer"   B = "is a smoker"

$\Pr[A|B] \approx 1.17 \times \Pr[A] \Rightarrow$ A, B positively correlated

<u>Note</u>: This <u>doesn't</u> necessarily imply that smoking <u>causes</u> lung cancer !

Independence

Defn: Events A, B are independent if

$$\Pr[A|B] = \Pr[A]$$

or equivalently if

$$\Pr[A \cap B] = \Pr[A] \times \Pr[B]$$

[Equivalent because $\Pr[A|B] = \dfrac{\Pr[A \cap B]}{\Pr[B]}$]

Independent or Not ?

1. 20 fair coin tosses

   A = all 20 tosses are H

   B = first 19 tosses are H

2. Roll 2 dice

   A = sum is ≥ 10

   B = first die shows 4

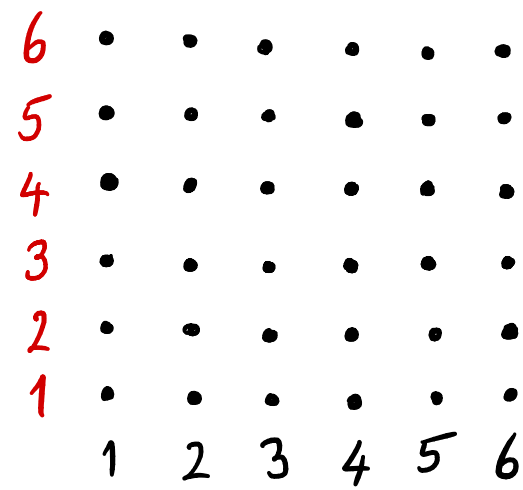

3. Toss 3 balls u.a.r. into 3 bins

   A = bin #1 is empty

   B = bin #2 is empty

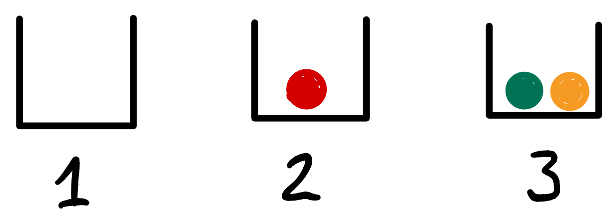

Mutual Independence

Defn: Events $A_1, \ldots, A_n$ are <u>mutually independent</u> if for all subsets $I \subseteq \{1, \ldots, n\}$

$$\Pr\left[\bigcap_{i \in I} A_i\right] = \prod_{i \in I} \Pr[A_i]$$

Example: 2 fair coin flips

A_1: "first flip is H" $\qquad \Pr[A_1] =$

A_2: "second flip is H" $\qquad \Pr[A_2] =$

A_3: "both flips the same (HH or TT)" $\qquad \Pr[A_3] =$

A_1, A_2: independent (obvious)

A_1, A_3: $\Pr[A_1 \cap A_3] = \Pr[HH] = 1/4 = \Pr[A_1]\Pr[A_3]$

A_2, A_3: same

BUT: $\Pr[A_1 \cap A_2 \cap A_3] =$

Independent Coin Flips

We often use independence to define prob. spaces

Example: Flipping a biased coin (Heads prob. p) twice

We want the flips to be independent, e.g.,

$$Pr[HT] = Pr[1st \text{ is } H] \times \underbrace{Pr[2nd \text{ is } H \mid 1st \text{ is } H]}_{= Pr[2nd \text{ is } H] \text{ (independence)}}$$

$$= p \times (1-p)$$

So we get

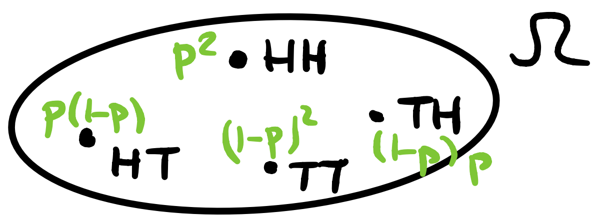

More generally, with n flips, for any seq. ω with i Heads and n-i Tails,

$$Pr[\omega] = p^i (1-p)^{n-i}$$

Intersections: Product Rule

Recall: $\Pr[A|B] = \dfrac{\Pr[A \cap B]}{\Pr[B]}$

This implies ...

Product Rule: For any events A, B

$$\Pr[A \cap B] = \Pr[A|B] \times \Pr[B] = \Pr[B|A] \times \Pr[A]$$

More generally ...

Product Rule: For any events $A_1, \ldots, A_n$

$$\Pr[A_1 \cap \ldots \cap A_n] = \Pr[A_1] \times \Pr[A_2|A_1] \times \ldots \times \Pr[A_n|A_1 \cap \ldots \cap A_{n-1}]$$

Product Rule: For any events $A_1, \dots, A_n$

$$\Pr[A_1 \cap \dots \cap A_n] = \Pr[A_1] \times \Pr[A_2 | A_1] \times \dots \times \Pr[A_n | A_1 \cap \dots \cap A_{n-1}]$$

Proof: By induction on n.

Base case $n=2$: basic product rule for 2 events

Inductive step $(n \geq 3)$:

$$\Pr[\underbrace{A_1 \cap \dots \cap A_{n-1}}_{B} \cap A_n] = \Pr[B] \times \Pr[A_n | B]$$

ind. hypothesis

$$= \Pr[A_1] \times \Pr[A_2 | A_1] \times \dots \times \Pr[A_{n-1} | A_1 \cap \dots \cap A_{n-2}]$$

$$\times \Pr[A_n | A_1 \cap \dots \cap A_{n-1}]$$

Unions of Events

Another dice game:

Roll two fair dice – you win if you roll at least one 6

Pr[roll 6 on one die] = 1/6

Pr[Win] = Pr[roll 6 on either die] = 1/6 + 1/6 = 1/3 ?

What if you roll 10 dice?

Pr[Win] = 1/6 + . . . + 1/6 = $\frac{10}{6}$???

Problem: You may roll more than one 6

Rolling 6's are not disjoint events

Thm: For any events A, B

$$\Pr[A\cup B] = \Pr[A] + \Pr[B] - \Pr[A\cap B]$$

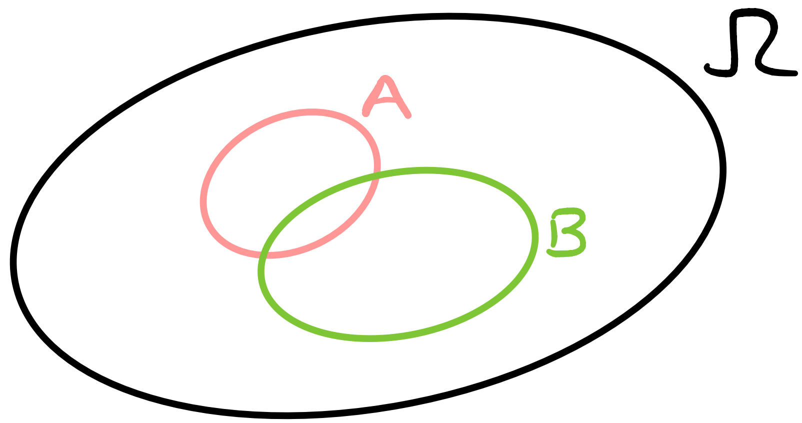

Proof: $\Pr[A\cup B] = \sum_{\omega\in A\cup B} \Pr[\omega]$

$$= \sum_{\omega\in A} \Pr[\omega] + \sum_{\omega\in B} \Pr[\omega] - \sum_{\omega\in A\cap B} \Pr[\omega]$$

$$= \Pr[A] + \Pr[B] - \Pr[A\cap B]$$

Note: If A, B are <u>disjoint</u> $(A\cap B=\emptyset)$ then $\Pr[A\cup B] = \Pr[A] + \Pr[B]$

Example:

Another dice game:

Roll two fair dice – you win if you roll at least one 6

$\Pr[\text{roll 6 on one die}] = 1/6$

A = roll 6 on first die

B = roll 6 on second die

$$\Pr[\text{Win}] = \Pr[A \cup B] = \Pr[A] + \Pr[B] - \Pr[A \cap B]$$

$$= 1/6 + 1/6 - 1/36$$

$$= 11/36$$

Inclusion-Exclusion

More generally, for any events $A_1, \ldots, A_n$

$$\Pr[A_1 \cup \ldots \cup A_n] = \sum_{i=1}^{n} \Pr[A_i] - \sum_{i<j} \Pr[A_i \cap A_j] + \sum_{i<j<k} \Pr[A_i \cap A_j \cap A_k] - \ldots \pm \Pr[A_1 \cap \ldots \cap A_n]$$

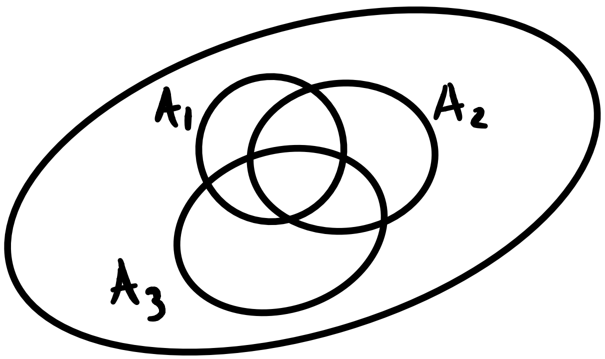

Proof: See inclusion-exclusion under "Counting"

Union Bound

Thm: For any events $A_1, \ldots, A_n$

$$\Pr[A_1 \cup \ldots \cup A_n] \leq \Pr[A_1] + \ldots + \Pr[A_n]$$

Proof: $\Pr[\bigcup_i A_i] = \sum_{\omega \in \bigcup A_i} \Pr[\omega]$

$$\leq \sum_{\omega \in A_1} \Pr[\omega] + \ldots + \sum_{\omega \in A_n} \Pr[\omega]$$

Later: We will see how useful this very simple upper bound can be!

Law of Total Probability

If $A_1, \dots, A_n$ are pairwise disjoint $(A_i \cap A_j = \emptyset \ \forall i \neq j)$ and $A_1 \cup \dots \cup A_n = \Omega$, then for any event B

$$\Pr[B] = \sum_{i=1}^{n} \Pr[B \cap A_i]$$

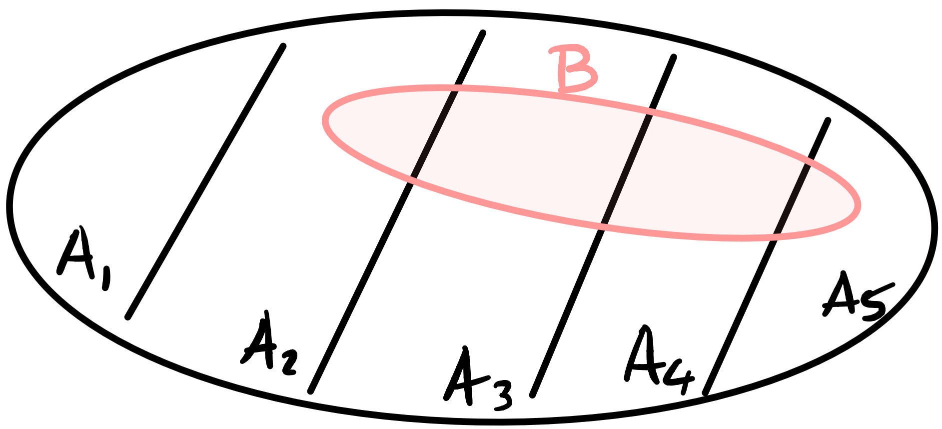

Proof: The events $B \cap A_i$ are pairwise disjoint and $B = \bigcup_i (B \cap A_i)$

Bayes Rule: For any events A, B with $\Pr[A] > 0$, $\Pr[B] > 0$, we have

$$\Pr[A|B] = \frac{\Pr[B|A]\Pr[A]}{\Pr[B]}$$

Proof: Statement is equivalent to

$$\Pr[A|B]\Pr[B] = \Pr[B|A]\Pr[A]$$

This is true because both sides $= \Pr[A \cap B]$

Bayes rule allows us to "flip the conditioning around", from $\Pr[B/A]$ to $\Pr[A|B]$

Example 1 : Two coins, Heads probs. $p = 1/2$ and $p = 3/5$

- pick a coin u.a.r. ("uniformly at random")
- flip the chosen coin

Suppose the flipped coin comes up Heads
What is the prob. we picked the biased coin ?

A = "picked biased coin"

B = "coin comes up Heads"

We know:

$$\Pr[A] = 1/2$$

$$\Pr[B|A] = 3/5 \qquad \Pr[B|\bar{A}] = 1/2$$

Goal : Compute $\Pr[A|B]$

A = "picked biased coin"

B = "coin comes up Heads"

We know: $\Pr[A] = 1/2$

$\Pr[B|A] = 3/5 \qquad \Pr[B|\bar{A}] = 1/2$

<u>Goal</u>: Compute $\Pr[A|B]$

<u>Bayes Rule</u>: $\Pr[A|B] = \dfrac{\Pr[B|A]\Pr[A]}{\Pr[B]} = \dfrac{3/5 \times 1/2}{\Pr[B]} = \dfrac{3/10}{\Pr[B]}$

What is $\Pr[B]$?

<u>Total Probability</u>: $\Pr[B] = \Pr[B|A]\Pr[A] + \Pr[B|\bar{A}]\Pr[\bar{A}]$

$= \left(\frac{3}{5} \times \frac{1}{2}\right) + \left(\frac{1}{2} \times \frac{1}{2}\right) = 11/20$

So $\Pr[A|B] = \dfrac{3/10}{11/20} = \boxed{6/11}$

Updated Bayes Rule

$$\Pr[A|B] = \frac{\Pr[B|A]\Pr[A]}{\underbrace{\Pr[B|A]\Pr[A] + \Pr[B|\bar{A}]\Pr[\bar{A}]}_{=\Pr[B]}}$$

More generally:

$$\Pr[A|B] = \frac{\Pr[B|A]\Pr[A]}{\underbrace{\sum_i \Pr[B|A_i]\Pr[A_i]}_{=\Pr[B]}}$$

where $A_1 \ldots A_n$ partitions Ω

E.g. 3 possible "Go" opponents, one chosen uniformly:

Opp. #1 wins w. prob. 90%
Opp. #2 - - - - - - - 60%
Opp. #3 - - - - - 20%

$$\Pr[\text{you lose}] = \left(\tfrac{1}{3} \times 0.9\right) + \left(\tfrac{1}{3} \times 0.6\right) + \left(\tfrac{1}{3} \times 0.2\right) \approx \boxed{0.57}$$

Example 2 : Medical Testing

Some disease affects 0.1% (=0.001) of population

A test has the following efficacy for a random person:

$$\left.\begin{aligned} \Pr[\text{test positive} \mid \text{sick}] &= 0.99 \\ \Pr[\text{test positive} \mid \text{not sick}] &= 0.01 \end{aligned}\right\}$$

false pos/neg rates are both 0.01

Q: A random person arrives & tests positive.
What is the likelihood this person is sick?

$\Pr[\text{pos.} \mid \text{sick}] = 0.99$

$\Pr[\text{pos.} \mid \text{not sick}] = 0.01$

$\Pr[\text{sick}] = 0.001$

Q: A random person arrives & tests positive. What is the likelihood this person is sick?

$$\Pr[\text{pos.} \mid \text{sick}] = 0.99 \qquad \Pr[\text{sick}] = 0.001$$

$$\Pr[\text{pos.} \mid \text{not sick}] = 0.01$$

<u>Bayes</u>:

$$\Pr[\text{sick} \mid \text{pos}] = \frac{\Pr[\text{pos} \mid \text{sick}]\,\Pr[\text{sick}]}{\Pr[\text{pos} \mid \text{sick}]\,\Pr[\text{sick}] + \Pr[\text{pos} \mid \text{not sick}]\,\Pr[\text{not sick}]}$$

$$= \frac{0.99 \times 0.001}{(0.99 \times 0.001) + (0.01 \times 0.999)}$$

$$\approx 0.09$$

Not a great test?

Reason: False pos. rate is large compared to % of sick people

Simpson's Paradox

On-time arrival performance of two airlines:

	Airline A			Airline B		
	#flights	#ontime	% ontime	#flights	#ontime	% ontime
L.A.	600	534	89%	200	188	94%
Chicago	250	176	70%	900	685	76%
Total	850	710	84%	1100	873	79%

Which airline would you fly
- into L.A. ?
- into Chicago ?
- overall ?

Explanation: Airline A has a much higher percentage of its flights into L.A., which has better performance than Chicago.

Math: Pick a random flight....

on Airline A

$\Pr[\text{on time} \mid LA] = 0.89$

$\Pr[\text{on time} \mid \text{Chicago}] = 0.70$

$\Pr[\text{on time}] = \Pr[\text{on time} \mid LA]\Pr[LA] + \Pr[\text{on time} \mid \text{Chic.}]\Pr[\text{Chic.}]$

$= \big((0.89 \times \Pr[LA]) + (0.70 \times \Pr[\text{Chic.}])\big)$

on Airline B

$\Pr[\text{on time} \mid LA] = 0.94$

$\Pr[\text{on time} \mid \text{Chic.}] = 0.76$

$\Pr[\text{on time}] = \ldots$

$= (0.94 \times \Pr[LA]) + (0.76 \times \Pr[\text{Chic}$

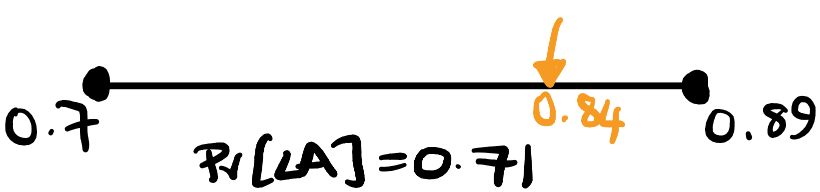

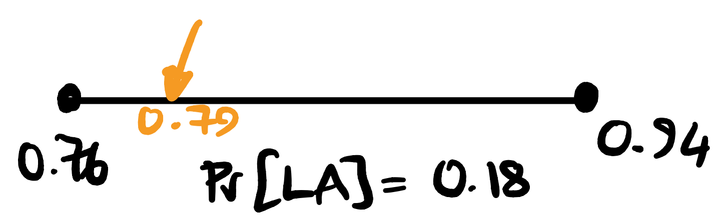

Summary

- Conditional probability
- Correlation & Independence
- Unions & intersections of events
- Bayes Rule & Total Probability Rule
- Inference; Simpson's Paradox