Uncertainty (Ch. 13)

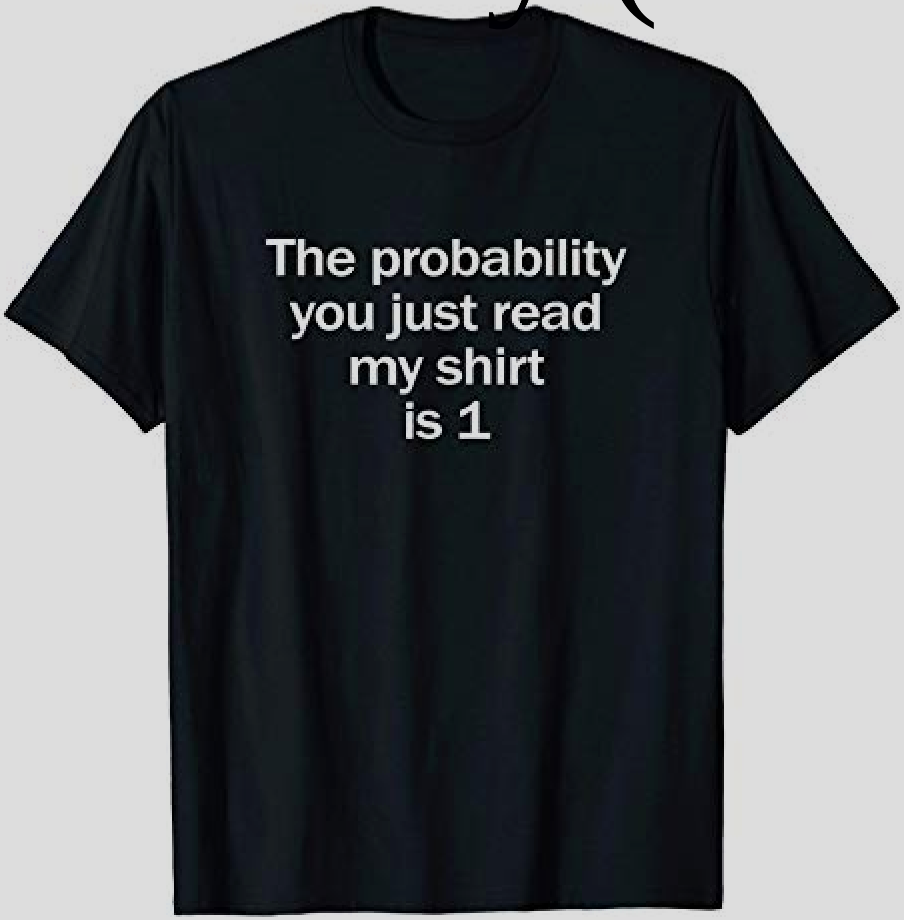

AlphaStar – Jan. 2019

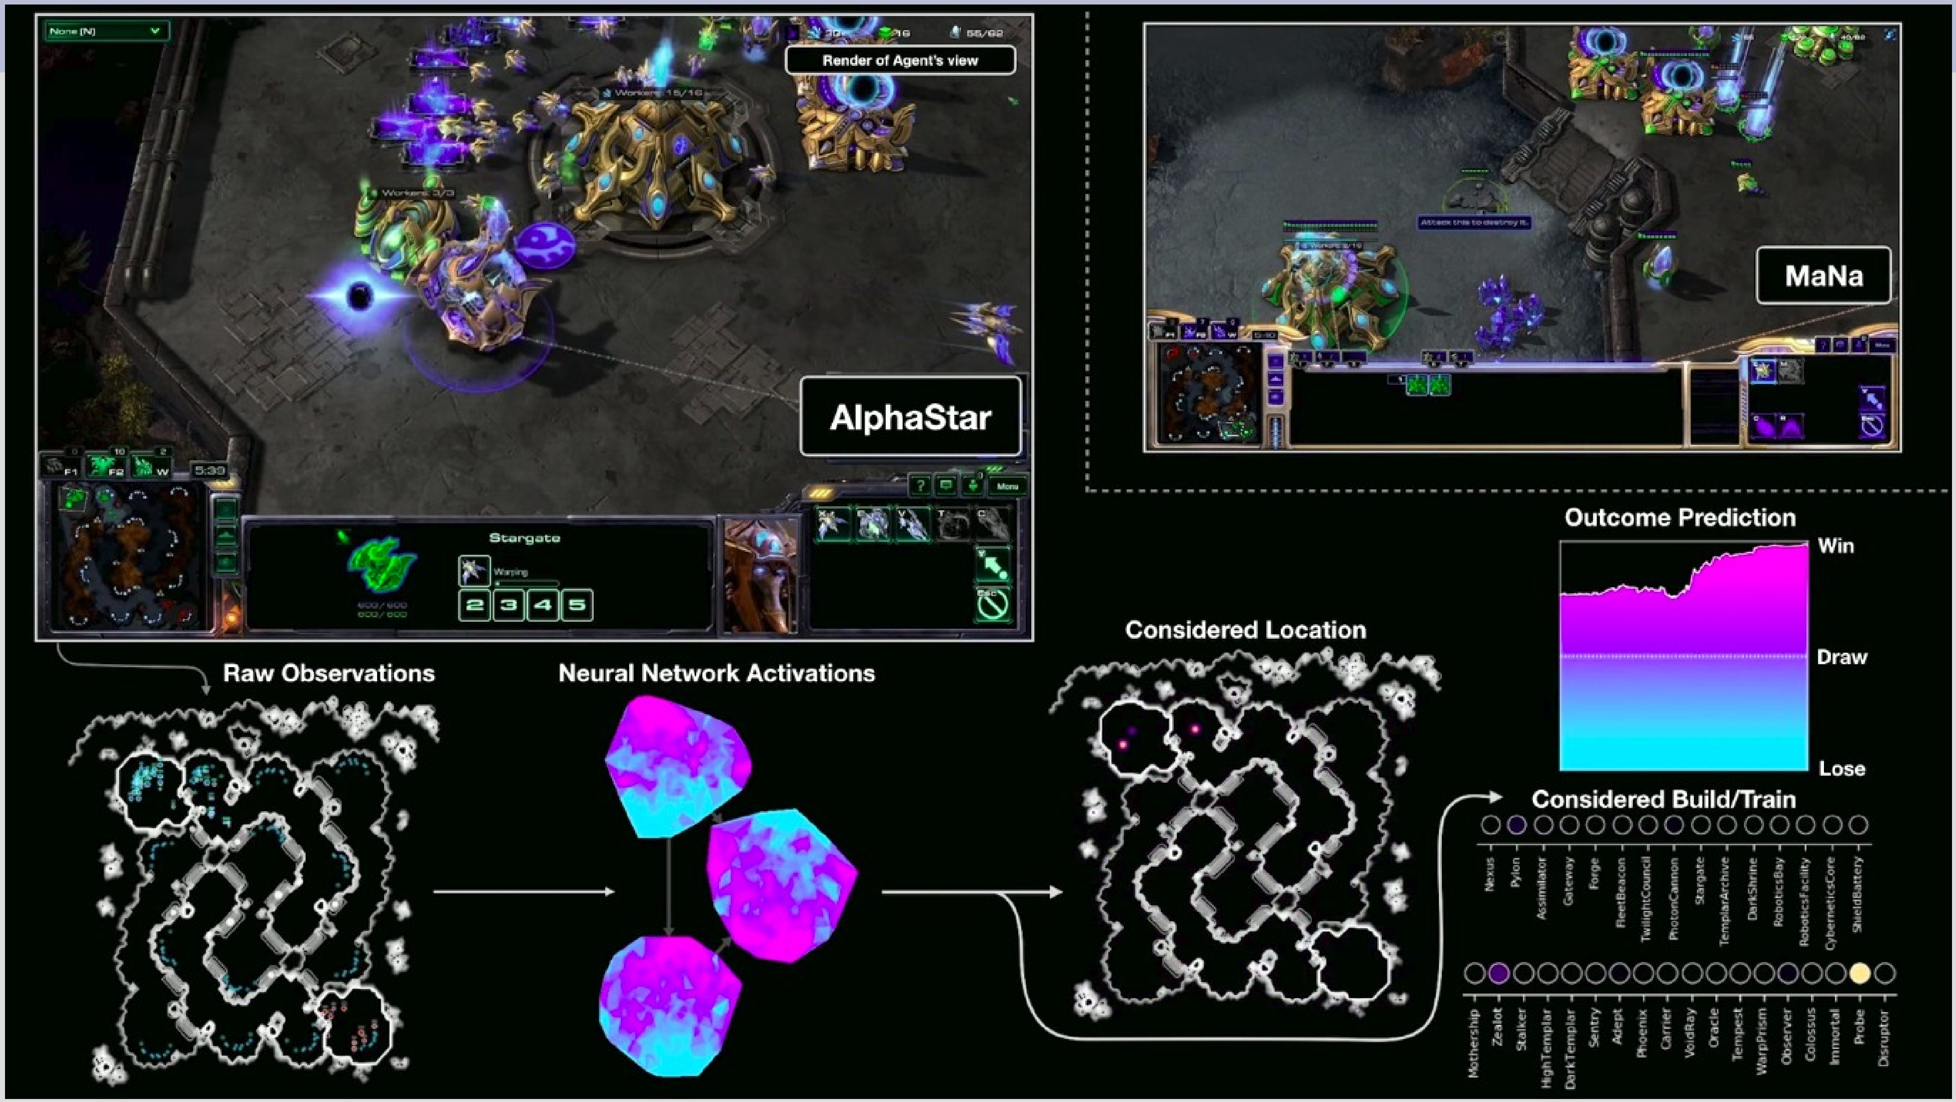

https://www.youtube.com/watch?v=cUTMhmVh1qs

Probability as values

Suppose we had a game where you payed $10 to play with the following situations:

1. You win $20 90% of the time, get $0 10%
2. Win $20 70%, get $0 30%
3. Win $20 0%, get $0 100%

Which games would you play?

For winning $20 or getting $0, what how low chance of winning before you should not play

Probability as values

Instead of paying $10 to a slot machine, you want to bet against another person

Again consider the following situations:

1. 20% win $5, 80% lose $5
2. 20% win $2, 80% lose $8
3. 20% win $8, 80% lose $2

If we assume the total “bet” is $10 (as in examples above)

Probability as values

Whats the (math) connection between paying $10 to a slot and betting $10 between people?

How would your strategy change if you bet $5 between people rather than betting $10?

Non-probability?

Consider the case:

P(A) = 0.2

P(B) = 0.3

P(A or B) = 0.1

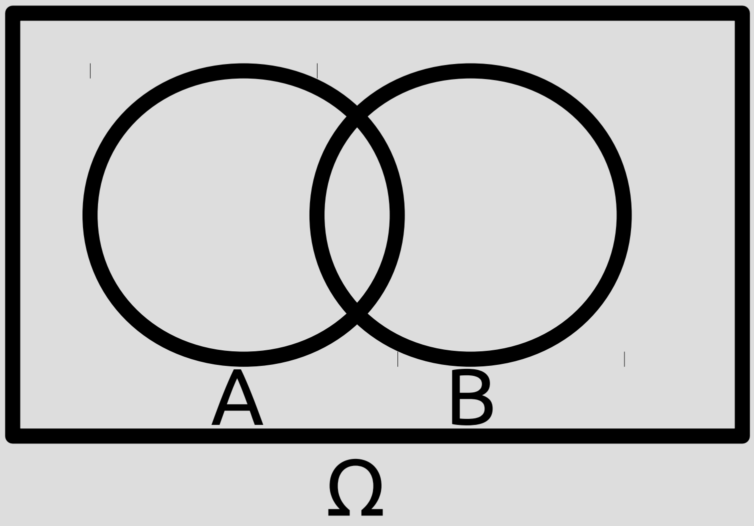

$$P(A \text{ or } B) = P(A) + P(B) - P(A, B)$$

Although this does not follow the rules of probability... would a robot that thinks this be in trouble?

Non-probability?

Consider the case:

P(A) = 0.2

P(B) = 0.3

P(A or B) = 0.1

	Robot bets	You bet
P(A)	2	8
P(B)	3	7
P(~(A or B))	16	4

Bet is neither (A or B) will happen

Yes! Assume you were betting against this robot and made the three bets above

The robot would think the first two fair and the last in their favor...

Non-probability?

If we look at the outcomes (regardless of what the probabilities are)

	Robot bets	You bet
P(A)	2	8
P(B)	3	7
P(~(A or B))	16	4

Robot gets:	A, B	~A, B	A, ~B	~A, ~B
Bet P(A)	+8	-2	+8	-2
Bet P(B)	+7	+7	-3	-3
Bet ~P(A or B)	-16	-16	-16	+4
Total	-1	-11	-11	-1

... no matter the outcome, robot will lose

Non-probability?

In fact, this is true for any “bad” set of probabilities

If you have non-mathematically sound probabilities, there is some betting strategy that will result in you always losing

This means someone could cheat our AI, so we will be careful to handle/use the rules of probability correctly

Prob. Manipulation

Common ways to manipulate probabilities:

1. Marginalize (i.e. sum out)

$$P(a) = \sum_{x \in B} P(a, x) = \sum_b P(a, b)$$

(informal)

2. Independence

$$P(A, B) = P(A)P(B)$$
$$P(A|B) = P(A)$$

(Note: upper case characters are the set for that variable, which will be True/False today)

Prob. Manipulation

Consider this table for P(breakfast,happy):

	happy	¬happy
breakfast	0.4	0.2
¬breakfast	0.1	0.3

$P(breakfast, \neg happy) = 0.2$

What is the chance that you are happy?

$$P(happy) = P(happy, breakfast) + P(happy, \neg breakfast)$$
$$= 0.4 + 0.1 = 0.5$$

Bayes' Rule

Bayes' Rule:

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

This equation isn't actually anything new, but this form of it is often used

This allows you to swap the direction of inference, which can be useful

Bayes' Rule

You went to the doctor and tested for a rare disease (1/1,000 people have it)

Test chances	Detected	Not Detected
Have disease	100%	0%
Just fine	1%	99%

If the test "detects" the disease,
what is the probability you are actually sick?

Bayes' Rule

d = disease, t = test detects

$$P(d) = 0.001, P(t|d) = 1, P(t|\neg d) = 0.01$$
$$P(\neg d) = 0.999, P(\neg t|d) = 0, P(\neg t|\neg d) = 0.99$$

Asked: P(d|t)=?

$$P(d|t) = \frac{P(t|d)P(d)}{P(t)} = \frac{(1)0.001}{P(t)}$$

How we get this?

Bayes' Rule

d = disease, t = test detects

$$P(d) = 0.001, P(t|d) = 1, P(t|\neg d) = 0.01$$
$$P(\neg d) = 0.999, P(\neg t|d) = 0, P(\neg t|\neg d) = 0.99$$

Asked: P(d|t)=?

$$P(d|t) = \frac{P(t|d)P(d)}{P(t)} = \frac{(1)0.001}{P(t)}$$

Sum over D in P(t,D)

Bayes' Rule

d = disease, t = test detects

$$P(d) = 0.001, P(t|d) = 1, P(t|\neg d) = 0.01$$
$$P(\neg d) = 0.999, P(\neg t|d) = 0, P(\neg t|\neg d) = 0.99$$

Asked: P(d|t)=?

$$P(d|t) = \frac{P(t|d)P(d)}{P(t)} = \frac{(1)0.001}{P(t)}$$
$$= \frac{(1)0.001}{P(t,d)+P(t,\neg d)} = \frac{(1)0.001}{P(t|d)P(d)+P(t|\neg d)P(\neg d)}$$
$$= \frac{(1)0.001}{(1)0.001+(0.01)0.999} = 0.090992$$

Bayes' Rule

I defined probabilities in terms of:

$P(d)$ and $P(t|d)$ (and negatives)

Is this a good choice?

What does the patient want to know?

Other ways could you store info?

Bayes' Rule

The patient probably wants to know:
$P(d)$ and $P(d|t)$

However, what is the most reliable (constant) probabilities here?

Bayes' Rule

The patient probably wants to know:

$P(d)$ and $P(d|t)$

However, what is the most reliable (constant) probabilities here?

$P(t|d)$

$P(d)$ and $P(d|t)$ both change if the rate of the disease goes up (both need re-estimate)

Bayes' Rule

There is another way to solve this problem without computing P(t)

$$P(d|t) = \frac{P(t|d)P(d)}{P(t)} = \frac{(1)0.001}{P(t)}$$

... and ...

$$P(\neg d|t) = \frac{P(t|\neg d)P(\neg d)}{P(t)} = \frac{(0.01)0.999}{P(t)}$$

Yet, we know: $P(d|t) + P(\neg d|t) = 1$
as you either have the disease or not (regardless of the test detection)

Bayes' Rule

This means you can treat P(t) as a constant and just normalize:

$$< P(d|t), P(\neg d|t) >=< \frac{(1)0.001}{P(t)}, \frac{(0.01)0.999}{P(t)} >$$
$$=< 0.001k, 0.00999k >$$

$$P(d|t) + P(\neg d|t) = 1 = 0.001k + 0.00999k, \text{ thus } k = 90.9008$$

$$< P(d|t), P(\neg d|t) >=< 0.001k, 0.00999k >=< 0.090992, 0.909008 >$$

Is this any different than the previous way?

Bayes' Rule

A small difference as you can use either:

$P(t)$ and $P(d,t)$ (or $P(t|d)P(d)$)

... or ...

$P(\neg d|t)$ and $P(d,t)$ (or $P(t|d)P(d)$)

... to get P(d|t)

Depending on how the data is stored, one way might be more efficient than the other

Independence

Independence between events can be thought of as the probabilities not effecting each other

$P(A|B) = P(A)$

A common example is flipping coins: if you have two coins, the outcome of the first coin flip does not effect the outcome of the second

This is quite often represented as:

$P(A, B) = P(A)P(B)$

Independence

Independence is useful as it can let you “drop” irrelevant information (easier computation)

In our disease/test problem we had:

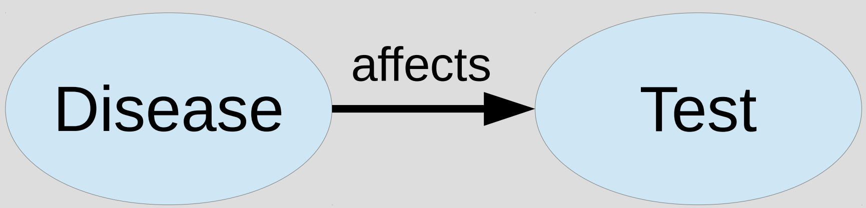

... so when computing P(d|t) we had to involve both variables (d ant t) to solve

Independence

If the test was independent of the disease (i.e. the doctor just flips a coin)

Then P(d|t) = P(d), which does not involve the right "Test" variable at all

(Note: often people think: $\text{independent} \Rightarrow P(A,B) = P(A)P(B)$ but actually also: $P(A,B) = P(A)P(B) \Rightarrow \text{independent}$)

Conditional Independence

If things are dependent, not all hope is lost

Consider a school example:

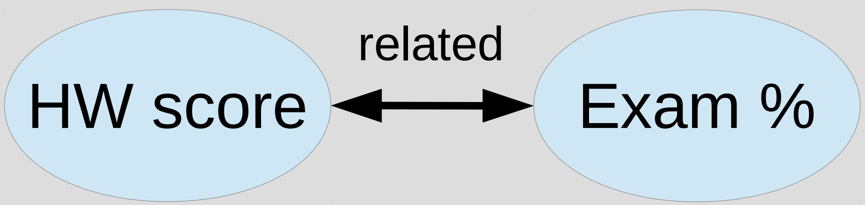

Hopefully your homework score should be some indication on how you do on an exam

Conditional Independence

But it would be more appropriate to probably have a third variable:

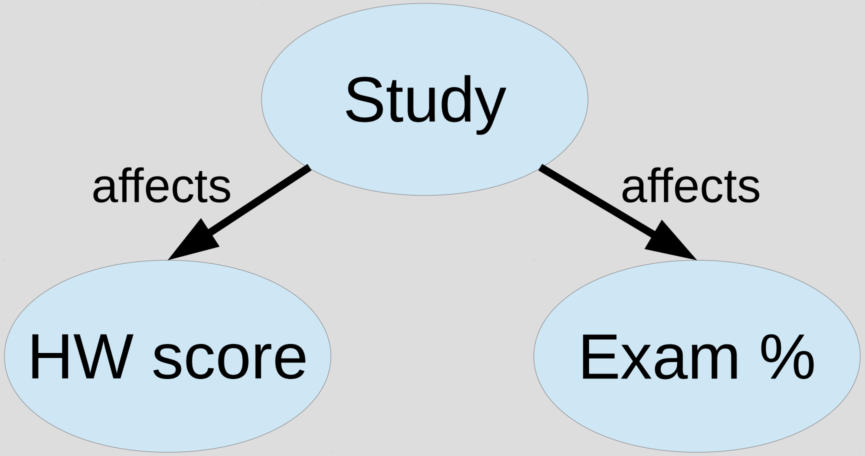

Here we would **not** say HW and Exam are independent, as $P(h, e) \neq P(h)P(e)$

However, as both homework and exams are driven by study...

Conditional Independence

We call this conditional independence, as

$P(h, e|s) = P(h|s)P(e|s)$

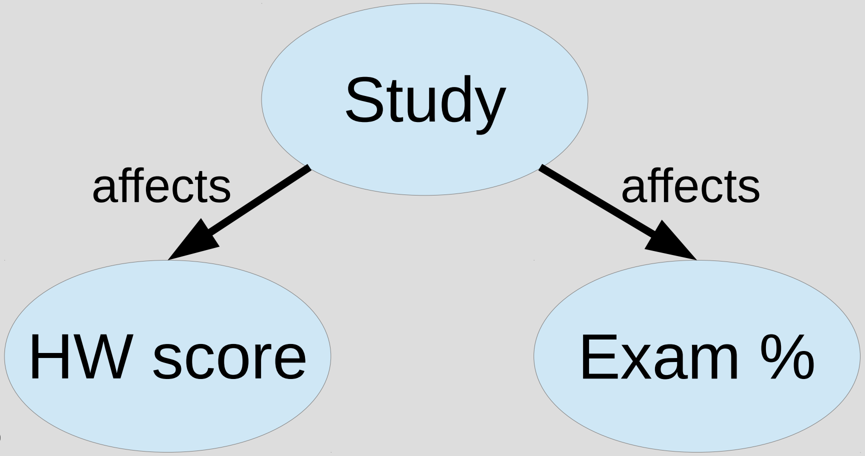

We will go over
how find if something
is conditionally independent next time

For now, the benefits are the same as normal
independence: less computation

General Independence

If we have n variables and want to store a table of $P(x_1, x_2, \ldots x_n)$, there would be 2^{n-1} entries in the table (assuming x is T/F)

Without independence, we would have to store $2^{n-1}-1$ of those (last skipped since sum to 1)

If all n variables are independent, you only need to store 1 value per variable

$O(n)$ is much better than $O(2^n)$

Prob. Manipulation

Recap: If you have $P(a,b|c,d)$

1. If you want to drop b (left of condition):

$$P(a|c,d) = \sum_b P(a,b|c,d)$$

2. If you want to drop c (right of condition):

   2.1. Hope conditionally independent:

$$P(a,b|d) = P(a,b|c,d)$$

   2.2. Else Method 1 + Conditional prob.

$$P(a,b|d) = \sum_c P(a,b,c|d) = \sum_c P(a,b|c,d)P(c|d)$$