

Probability and Decision Theory

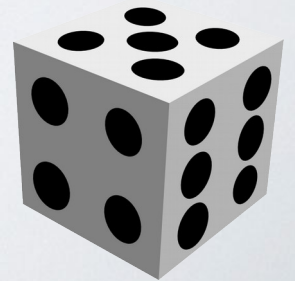
Robert Platt
Northeastern University

Some images and slides are used from:

1. AIMA
2. Chris Amato
3. Stacy Marsella

QUANTIFYING UNCERTAINTY WITH PROBABILITIES

- Generally we may use probabilities due to
 - Ignorance
 - E.g., we don't know our opponent
 - Laziness
 - It is too difficult or takes too much effort to model the event in detail
- The event is inherently random



HOW DO WE INTERPRET PROBABILITIES

- Frequentist view
 - Objective: If I throw a die many times it will come up 3 one sixth of the times
- Strength of belief view (Bayesian)
 - Subjective: How strongly do I, should I, believe it will come up 3
 - These are not claims of a “probabilistic tendency” in the current situation (but might be learned from past experience)
 - Probabilities of propositions change with new evidence:
 - e.g., $P(\text{get to airport on time} | \text{no reported accidents, 5 a.m.}) = 0.15$

(Discrete) Random variables

What is a random variable?

Suppose that the variable a denotes the outcome of a role of a single six-sided die:

$$a \in \{1, 2, 3, 4, 5, 6\} = A$$

a is a random variable

this is the *domain* of a

Another example:

Suppose b denotes whether it is raining or clear outside:

$$b \in \{rain, clear\} = B$$

Probability distribution

A probability distribution associates each with a probability of occurrence, represented by a *probability mass function (pmf)*.

A probability table is one way to encode the distribution:

$$a \in \{1, 2, 3, 4, 5, 6\} = A \quad b \in \{rain, clear\} = B$$

a	P(a)
1	1/6
2	1/6
3	1/6
4	1/6
5	1/6
6	1/6

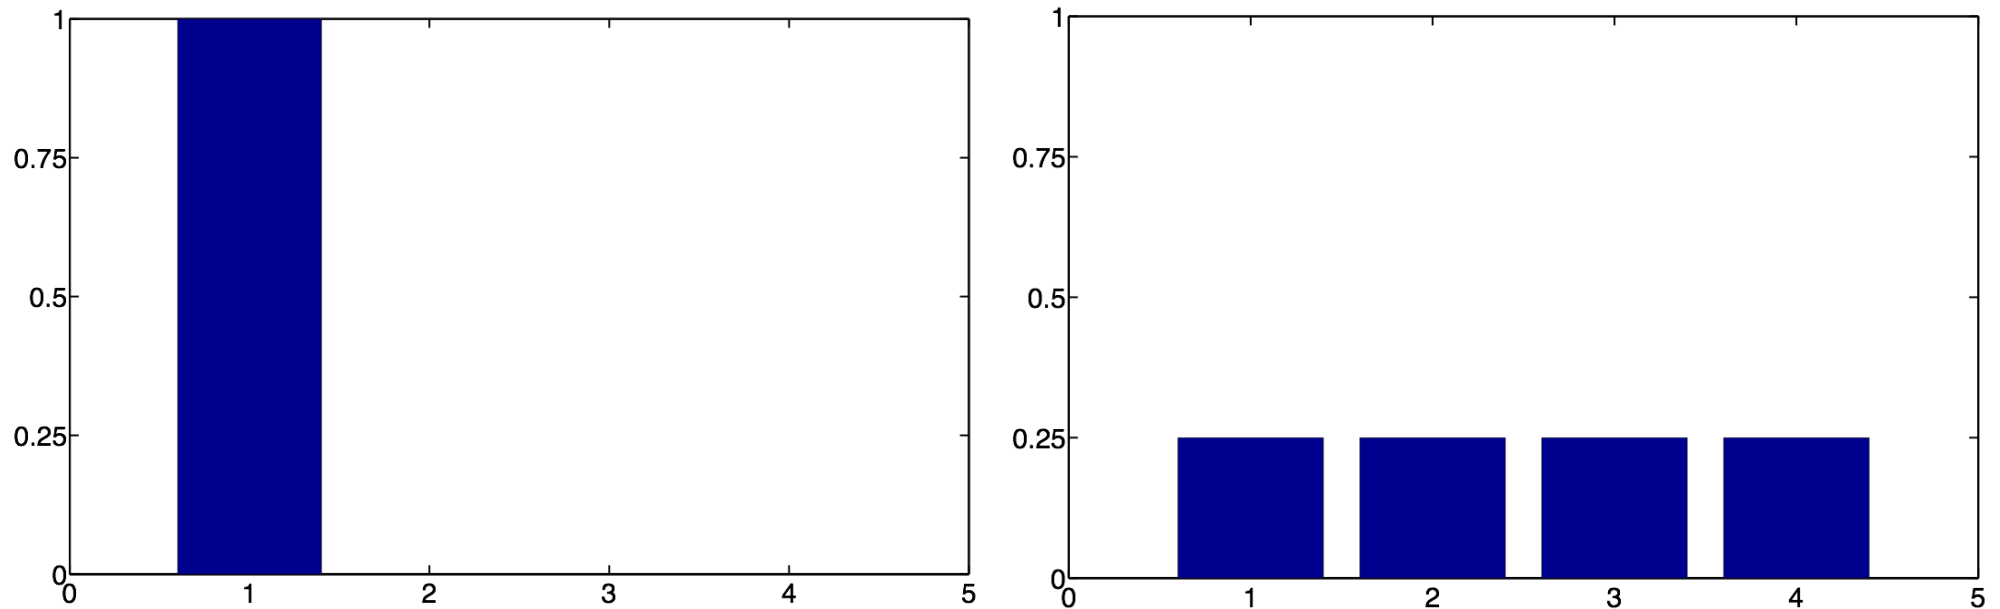
b	P(b)
rain	1/4
clear	3/4

All probability distributions must satisfy the following:

1. $\forall a \in A, a \geq 0$

2. $\sum_{a \in A} a = 1$

Example pmfs



Two pmfs over a state space of $X = \{1, 2, 3, 4\}$

Writing probabilities

a	P(a)
1	1/6
2	1/6
3	1/6
4	1/6
5	1/6
6	1/6

b	P(b)
rain	1/4
clear	3/4

For example: $p(a = 2) = 1/6$
 $p(b = \text{clear}) = 3/4$

But, sometimes we will abbreviate this as: $p(2) = 1/6$

$$p(\text{clear}) = 3/4$$

Types of random variables

Propositional or Boolean random variables

- e.g., *Cavity* (do I have a cavity?)
- *Cavity = true* is a proposition, also written *cavity*

Discrete random variables (finite or infinite)

- e.g., *Weather* is one of $\langle \textit{sunny}, \textit{rain}, \textit{cloudy}, \textit{snow} \rangle$
- *Weather = rain* is a proposition
- Values must be exhaustive and mutually exclusive

Continuous random variables (bounded or unbounded)

- e.g., $\text{Temp} < 22.0$

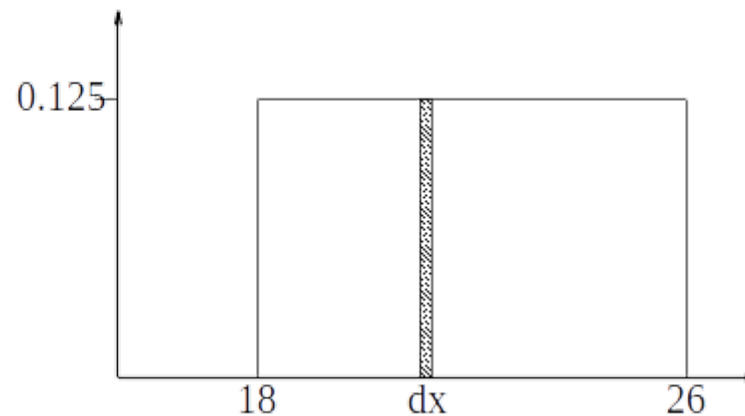
Continuous random variables

Cumulate distribution function (cdf), $F(q)=(X<q)$ with $P(a<X\leq b)=F(b)-F(a)$

Probability density function (pdf), $f(x)=\frac{d}{dx}F(x)$ with $P(a<X\leq b)=\int_a^b f(x)$

Express distribution as a parameterized function of value:

- e.g., $P(X = x) = U[18, 26](x) = \text{uniform density between 18 and 26}$



Here P is a density; integrates to 1.

$P(X = 20.5) = 0.125$ really means $\lim_{dx \rightarrow 0} P(20.5 \leq X \leq 20.5 + dx) / dx = 0.125$

Joint probability distributions

Given random variables: X_1, X_2, \dots, X_n

The *joint distribution* is a probability assignment to all combinations: $P(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$

or:
$$P(x_1, x_2, \dots, x_n)$$

Sometimes written as:
$$P(X_1 = x_1 \wedge X_2 = x_2 \wedge \dots \wedge X_n = x_n)$$

As with single-variate distributions, joint distributions must satisfy:

1.
$$P(x_1, x_2, \dots, x_n) \geq 0$$

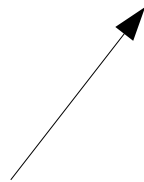
2.
$$\sum_{x_1, \dots, x_n} P(x_1, x_2, \dots, x_n) = 1$$

Prior or unconditional probabilities of propositions
e.g., $P(\text{Cavity} = \text{true}) = 0.1$ and $P(\text{Weather} = \text{sunny}) = 0.72$
correspond to belief prior to arrival of any (new) evidence

Joint probability distributions

Joint distributions are typically written in table form:

T	W	$P(T,W)$
Warm	snow	0.1
Warm	hail	0.3
Cold	snow	0.5
Cold	hail	0.1

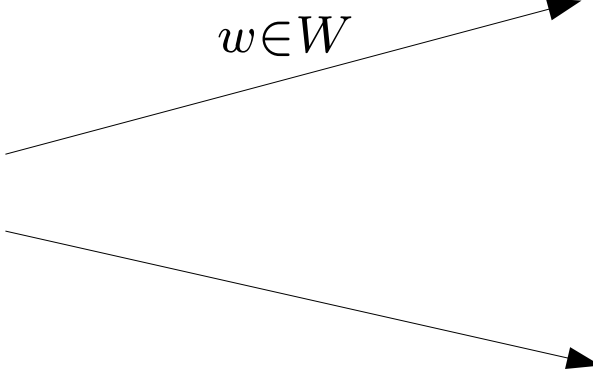


An event is a set of particular outcomes
– e.g. it's warm and hailing...

Marginalization

Given $P(T,W)$, calculate $P(T)$ or $P(W)$...

T	W	P(T,W)
Warm	snow	0.1
Warm	hail	0.3
Cold	snow	0.4
Cold	hail	0.2

$$P(T) = \sum_{w \in W} P(T, w)$$


T	P(T)
Warm	0.4
Cold	0.6

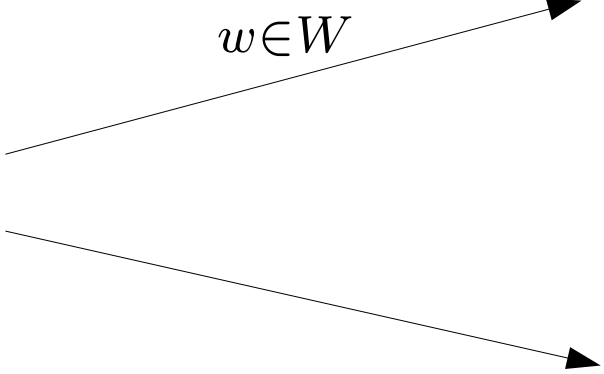
$$P(W) = \sum_{t \in T} P(t, W)$$

W	P(W)
snow	0.5
hail	0.5

Marginalization

Given $P(T,W)$, calculate $P(T)$ or $P(W)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3

$$P(T) = \sum_{w \in W} P(T, w)$$


T	P(T)
Warm	?
Cold	?

$$P(W) = \sum_{t \in T} P(t, W)$$

W	P(W)
snow	?
hail	?

Conditional Probability

At a regular checkup, assume the unconditional probability that you have a cavity is 0.2

But what if I have a toothache? Does this change the probability of a cavity?

$$P(\text{cavity} \mid \text{toothache})$$

Probability of a cavity GIVEN a toothache

Conditional Probability

At a regular checkup, assume the unconditional probability that you have a cavity is 0.2

But what if I have a toothache? Does this change the probability of a cavity?

$$P(\text{cavity} \mid \text{toothache})$$

Probability of a cavity GIVEN a toothache

Definition of conditional probability:
$$P(a|b) = \frac{P(a, b)}{P(b)}$$

Can be rewritten as the Product Rule:
$$P(a, b) = P(a|b)P(b) = P(b|a)P(a)$$

$$P(x_n, \dots, x_1) = P(x_n | x_{n-1}, \dots, x_1) P(x_{n-1}, \dots, x_1)$$

The Chain Rule

The product rule can be applied iteratively to obtain the chain rule:

$$P(x_n, \dots, x_1) = P(x_n | x_{n-1}, \dots, x_1) P(x_{n-1}, \dots, x_1)$$
$$= P(x_n | x_{n-1}, \dots, x_1) P(x_{n-1} | x_{n-2}, \dots, x_1) \dots P(x_1)$$

Chain rule



Infer Conditional Probabilities

	<i>toothache</i>		\neg <i>toothache</i>	
	<i>catch</i>	\neg <i>catch</i>	<i>catch</i>	\neg <i>catch</i>
<i>cavity</i>	0.108	0.012	0.072	0.008
\neg <i>cavity</i>	0.016	0.064	0.144	0.576

We want to calculate the probability of a cavity given that we have a toothache

Recall that: $P(a|b) = \frac{P(a, b)}{P(b)}$

$$P(\text{cavity} \mid \text{toothache}) = \frac{P(\text{cavity} \wedge \text{toothache})}{P(\text{toothache})} = \frac{0.108 + 0.012}{0.108 + 0.012 + 0.016 + 0.064} = 0.6$$

$$P(\neg \text{cavity} \mid \text{toothache}) = \frac{P(\neg \text{cavity} \wedge \text{toothache})}{P(\text{toothache})} = ???$$

Infer Conditional Probabilities

$P(\text{snow}|\text{warm})$ = Probability that it will snow *given* that it is warm

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3




W	$P(W T=\text{warm})$
snow	?
hail	?

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3




W	$P(W T=\text{warm})$
snow	?
hail	?

$$P(\text{snow}|\text{warm}) = \frac{P(\text{warm}, \text{snow})}{P(\text{warm})} = \frac{P(\text{warm}, \text{snow})}{P(\text{warm}, \text{hail}) + P(\text{warm}, \text{snow})}$$

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3




W	$P(W T=\text{warm})$
snow	0.6
hail	?

$$\begin{aligned} P(\text{snow}|\text{warm}) &= \frac{P(\text{warm}, \text{snow})}{P(\text{warm})} = \frac{P(\text{warm}, \text{snow})}{P(\text{warm}, \text{hail}) + P(\text{warm}, \text{snow})} \\ &= \frac{0.3}{0.2 + 0.3} \end{aligned}$$

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3




W	$P(W T=\text{warm})$
snow	0.6
hail	?

$$\begin{aligned} P(\text{snow}|\text{warm}) &= \frac{P(\text{warm}, \text{snow})}{P(\text{warm})} = \frac{P(\text{warm}, \text{snow})}{P(\text{warm}, \text{hail}) + P(\text{warm}, \text{snow})} \\ &= \frac{0.3}{0.2 + 0.3} \end{aligned}$$

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3



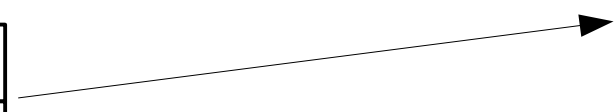
W	$P(W T=\text{warm})$
snow	0.6
hail	0.4

$$\begin{aligned} P(\text{snow}|\text{warm}) &= \frac{P(\text{warm}, \text{snow})}{P(\text{warm})} = \frac{P(\text{warm}, \text{snow})}{P(\text{warm}, \text{hail}) + P(\text{warm}, \text{snow})} \\ &= \frac{0.3}{0.2 + 0.3} \end{aligned}$$

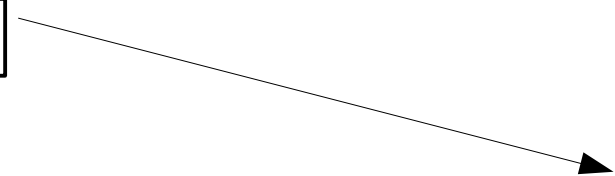
Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3



W	$P(W T=\text{warm})$
snow	0.6
hail	0.4

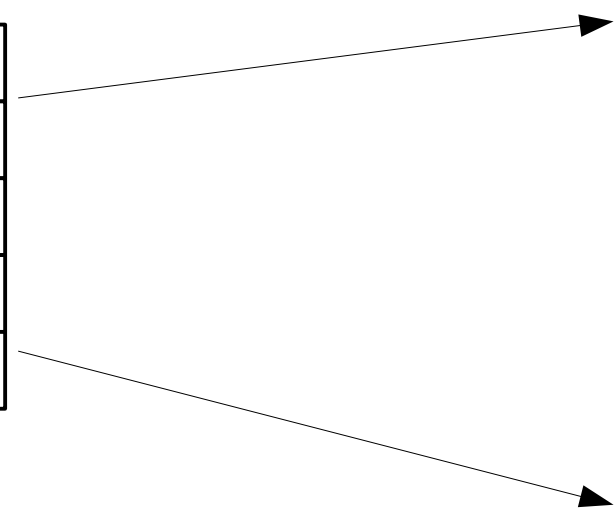


W	$P(W T=\text{cold})$
snow	?
hail	?

Infer Conditional Probabilities

Given $P(T,W)$, calculate $P(T|w)$ or $P(W|t)$...

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3



W	$P(W T=\text{warm})$
snow	0.6
hail	0.4

W	$P(W T=\text{cold})$
snow	0.4
hail	0.6

Normalization

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3



W	P(W T=warm)
snow	0.6
hail	0.4

Can we avoid explicitly computing this denominator?

$$P(\text{snow}|\text{warm}) = \frac{P(\text{warm}, \text{snow})}{P(\text{warm}, \text{hail}) + P(\text{warm}, \text{snow})}$$

Any ideas?

Normalization

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.2
Cold	snow	0.2
Cold	hail	0.3

W	P(W,T=warm)
snow	0.3
hail	0.2

$$P(W|t) = \frac{P(W,t)}{P(t)}$$

Two steps:

1. Copy entries

2. Scale them up so
that entries sum to 1

W	P(W T=warm)
snow	0.6
hail	0.4

W	P(W T=warm)
snow	0.6
hail	0.4

Normalization

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.4
Cold	snow	0.2
Cold	hail	0.1

Two steps:

1. Copy entries

T	P(T,W=hail)
warm	?
cold	?

2. Scale them up so
that entries sum to 1

T	P(T W=hail)
warm	?
cold	?

Normalization

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.4
Cold	snow	0.2
Cold	hail	0.1

Two steps:

1. Copy entries

T	P(T,W=hail)
warm	0.4
cold	0.1

2. Scale them up so
that entries sum to 1

T	P(T W=hail)
warm	?
cold	?

Normalization

T	W	P(T,W)
Warm	snow	0.3
Warm	hail	0.4
Cold	snow	0.2
Cold	hail	0.1

Two steps:

1. Copy entries

T	P(T,W=hail)
warm	0.4
cold	0.1

2. Scale them up so
that entries sum to 1

T	P(T W=hail)
warm	0.8
cold	0.2

$$P(W|t) = \frac{P(W, t)}{P(t)}$$

The only purpose of this denominator is to make the distribution sum to one.
– we achieve the same thing by scaling.

A Problem...

So far we have done inference using the full joint distribution

But, this doesn't scale well with complex domains...

- e.g. 30 boolean variables requires us to specify 2^{30} probabilities!

Way too many!



Possible solution: exploit independence!

INDEPENDENCE

- Expand our dentistry joint distribution by adding a fourth variable, weather
 - $P(\text{Toothache}, \text{Catch}, \text{Cavity}, \text{Weather})$
- *Now have a table of 32 entries*
 - (Assuming $\text{Weather} = \{\text{cloudy}, \text{rain}, \text{sunny}, \text{snow}\}$)
- But consider the relation these variables have:
 - $P(\text{toothache}, \text{catch}, \text{cavity}, \text{cloudy})$
 - *Does cloudy influence dentistry?*

INDEPENDENCE

- By the product rule:
 - $P(\text{toothache}, \text{catch}, \text{cavity}, \text{cloudy}) = P(\text{cloudy} | \text{toothache}, \text{catch}, \text{cavity}) \times P(\text{toothache}, \text{catch}, \text{cavity})$
- Reasonable to assume dental facts and the weather don't influence each other, so:
 - $P(\text{cloudy} | \text{toothache}, \text{catch}, \text{cavity}) = P(\text{cloudy})$
- Therefore: $P(\text{toothache}, \text{catch}, \text{cavity}, \text{cloudy}) = P(\text{cloudy}) \times P(\text{toothache}, \text{catch}, \text{cavity})$
- Reducing the 32 entry table to an 8 element table and 4 element table
 - Calculation of joint is a product of the entries in these tables

INDEPENDENCE

- *A and B are independent iff $P(X, Y) = P(X)P(Y)$*
 - *$P(\text{Toothache}, \text{Catch}, \text{Cavity}, \text{Weather}) = P(\text{Toothache}, \text{Catch}, \text{Cavity})P(\text{Weather})$*
- *In other words the joint distribution factors into a product of two smaller, simpler distributions*
- *Equivalently*
 - *$P(X|Y) = P(X)$ or $P(Y|X) = P(Y)$*
- *If two events are independent, knowing one has occurred does not give us any new information about the other*

INDEPENDENT?

- Formally independence is

- $P(X|Y) = P(X)$ or
- $P(Y|X) = P(Y)$ or
- $P(X, Y) = P(X)P(Y)$

$$P_1(T, W)$$

T	W	P
hot	sun	0.4
hot	rain	0.1
cold	sun	0.2
cold	rain	0.3

$$P(T)$$

T	P
hot	0.5
cold	0.5

$$P(W)$$

W	P
sun	0.6
rain	0.4

$$P_2(T, W)$$

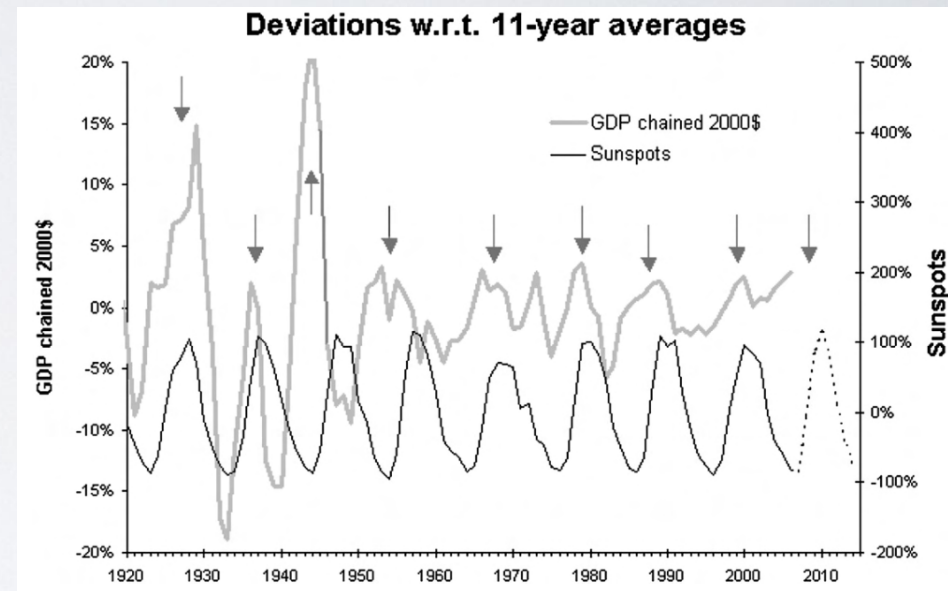
T	W	P
hot	sun	0.3
hot	rain	0.2
cold	sun	0.3
cold	rain	0.2

INDEPENDENCE

- Formally independence is (equivalently)
 - *In terms of conditional distributions*
 - $P(X|Y) = P(X)$ or $P(Y|X) = P(Y)$
 - *In terms of joint distributions*
 - $P(X, Y) = P(X)P(Y)$
- Independence can greatly reduce the probabilities that need to be specified

INDEPENDENCE

- However full independence is rare
- Just as separate subgraphs were rare in CSPs
- In empirical data especially, it's rare not to see some interaction (real or not)
- Sunspot activity and GDP,
- Sunspot activity and stock market
- Sunspot activity and ...
- We need to consider a more limited form of independence



CONDITIONAL INDEPENDENCE

We will consider a more limited form of independence

Conditional Independence:

$$P(X, Y | Z) = P(X | Z)P(Y | Z)$$

$P(\text{Toothache}, \text{Cavity}, \text{Catch})$ has $8 - 1 = 7$ independent entries

If I have a cavity, the probability that the probe catches in it doesn't depend on whether I have a toothache:

$$P(\text{catch} | \text{toothache}, \text{cavity}) = P(\text{catch} | \text{cavity})$$

The same independence holds if I haven't got a cavity:

$$P(\text{catch} | \text{toothache}, \neg \text{cavity}) = P(\text{catch} | \neg \text{cavity})$$

Catch is conditionally independent of Toothache given Cavity:

$$P(\text{Catch} | \text{Toothache}, \text{Cavity}) = P(\text{Catch} | \text{Cavity})$$

Equivalent statements:

$$P(\text{Toothache} | \text{Catch}, \text{Cavity}) = P(\text{Toothache} | \text{Cavity})$$

$$P(\text{Toothache}, \text{Catch} | \text{Cavity}) = P(\text{Toothache} | \text{Cavity})P(\text{Catch} | \text{Cavity})$$

CONDITIONAL INDEPENDENCE

- Write out full joint distribution using chain rule:
- $P(\textit{Toothache}, \textit{Catch}, \textit{Cavity})$
 - $= P(\textit{Toothache}|\textit{Catch}, \textit{Cavity})P(\textit{Catch}, \textit{Cavity})$
 - $= P(\textit{Toothache}|\textit{Catch}, \textit{Cavity})P(\textit{Catch}|\textit{Cavity})P(\textit{Cavity})$
 - $= P(\textit{Toothache}|\textit{Cavity})P(\textit{Catch}|\textit{Cavity})P(\textit{Cavity})$
- $2 + 2 + 1 = 5$ independent numbers
- In most cases, the use of conditional independence reduces the size of the representation of the joint distribution from exponential in n to linear in n
- Conditional independence is our most basic and robust form of knowledge about uncertain environments

Bayes Rule

$$P(a|b) = \frac{P(b|a)P(a)}{P(b)}$$



Thomas Bayes (1701 – 1761):

- English statistician, philosopher and Presbyterian minister
- formulated a specific case of the formula above
- his work later published/generalized by Richard Price

Bayes Rule

$$P(a|b) = \frac{P(b|a)P(a)}{P(b)}$$

It's easy to derive from the product rule:

$$P(a, b) = P(b|a)P(a) = \underbrace{P(a|b)P(b)}$$



Solve for this

Using Bayes Rule

$$P(a|b) = \frac{P(b|a)P(a)}{P(b)}$$



$$P(cause|effect) = \frac{P(effect|cause)P(cause)}{P(effect)}$$

Using Bayes Rule

$$P(a|b) = \frac{P(b|a)P(a)}{P(b)}$$



$$P(\textit{cause}|\textit{effect}) = \frac{P(\textit{effect}|\textit{cause})P(\textit{cause})}{P(\textit{effect})}$$

But harder to estimate this

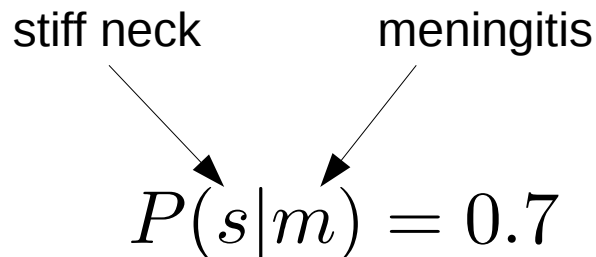
It's often easier to estimate this

Bayes Rule Example

$$P(\textit{cause}|\textit{effect}) = \frac{P(\textit{effect}|\textit{cause})P(\textit{cause})}{P(\textit{effect})}$$

Suppose you have a stiff neck...

Suppose there is a 70% chance of meningitis if you have a stiff neck:



stiff neck meningitis

$$P(s|m) = 0.7$$

What are the chances that you have meningitis?

Bayes Rule Example

$$P(\textit{cause}|\textit{effect}) = \frac{P(\textit{effect}|\textit{cause})P(\textit{cause})}{P(\textit{effect})}$$

Suppose you have a stiff neck...

Suppose there is a 70% chance of meningitis if you have a stiff neck:

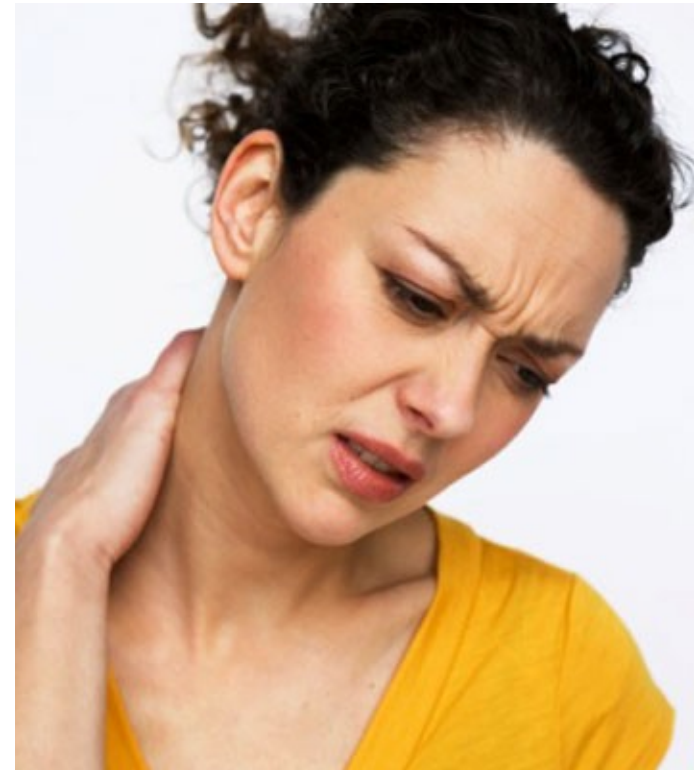
stiff neck meningitis

↘ ↘

$$P(s|m) = 0.7$$

What are the chances that you have meningitis?

We need a little more information...



Bayes Rule Example

$$P(\text{cause}|\text{effect}) = \frac{P(\text{effect}|\text{cause})P(\text{cause})}{P(\text{effect})}$$

$$P(s|m) = 0.7$$

$$P(s) = 0.01$$



Prior probability of stiff neck

$$P(m) = \frac{1}{50000}$$



Prior probability of meningitis

$$P(m|s) = \frac{P(s|m)P(m)}{P(s)} = \frac{0.7 \times \frac{1}{50000}}{0.01} = 0.0014$$

Bayes Rule Example

Given:

W	P(W)
snow	0.8
hail	0.2

T	W	P(T W)
Warm	snow	0.3
Warm	hail	0.4
Cold	snow	0.7
Cold	hail	0.6

Calculate P(W|warm):

$$P(W|warm) = \frac{P(warm|W)P(W)}{P(warm)}$$

Bayes Rule Example

Given:

W	P(W)
snow	0.8
hail	0.2

T	W	P(T W)
Warm	snow	0.3
Warm	hail	0.4
Cold	snow	0.7
Cold	hail	0.6

Calculate P(W|warm):

$$P(W|warm) = \frac{P(warm|W)P(W)}{P(warm)}$$

$$P(hail|warm) = \frac{0.4 \times 0.2}{P(warm)} = \frac{0.08}{P(warm)} = 0.25$$

$$P(snow|warm) = \frac{0.3 \times 0.8}{P(warm)} = \frac{0.24}{P(warm)} = 0.75$$

normalize

SO FAR: IMPORTANT FORMULAS

- *Basic axioms*
 - $P(a) \geq 0$
 - $\sum_{a \in X} P(a) = 1$
 - $P(\neg a) = 1 - P(a)$
 - (holds for joint distributions as well)
- *Inclusion-exclusion principle*
 - $P(a \vee b) = P(a) + P(b) - P(a \wedge b)$
- *Marginal probability*
 - $P(Y) = \sum_{z \in Z} P(Y, z)$
- *X and Y are independent **IF***
 - $P(X, Y) = P(X)P(Y)$
 - $P(X|Y) = P(X)$ or $P(Y|X) = P(Y)$ or
- *X, Y are conditional independent given Z **IF***
 - $P(X, Y|Z) = P(X|Z) P(Y|Z)$
 - $P(X|Y, Z) = P(X|Z)$
 - $P(Y|X, Z) = P(Y|Z)$
- *Definition of conditional probability:*
 - $P(a | b) = \frac{P(a \wedge b)}{P(b)}$
 - $P(X | Y) = \frac{P(X, Y)}{P(Y)}$
- *Product Rule:*
 - $P(a \wedge b) = P(a|b) P(b)$ or $P(a \wedge b) = P(b|a) P(a)$
 - $P(X_n, \dots, X_1) = P(X_n | X_{n-1}, \dots, X_1) P(X_{n-1}, \dots, X_1)$
- *Chain Rule:*
 - $P(X_n, \dots, X_1) = P(X_n | X_{n-1}, \dots, X_1) P(X_{n-1} | X_{n-2}, \dots, X_1) \dots P(X_1)$
- *Bayes Rule*
 - $P(Y|X) = \frac{P(X|Y) P(Y)}{P(X)}$
 - $P(Y|X, e) = \frac{P(X|Y, e) P(Y | e)}{P(X | e)}$

MAKING DECISIONS UNDER UNCERTAINTY

- Going to the airport: A_x means leaving X minutes before the flight
- Suppose I believe the following:
 - $P(A_{25} \text{ gets me there on time} | \dots) = 0.04$
 - $P(A_{90} \text{ gets me there on time} | \dots) = 0.70$
 - $P(A_{120} \text{ gets me there on time} | \dots) = 0.95$
 - $P(A_{1440} \text{ gets me there on time} | \dots) = 0.9999$
- Which action to choose?

MAKING DECISIONS UNDER UNCERTAINTY

- Suppose I believe the following:
 - $P(A_{25} \text{ gets me there on time} | \dots) = 0.04$
 - $P(A_{90} \text{ gets me there on time} | \dots) = 0.70$
 - $P(A_{120} \text{ gets me there on time} | \dots) = 0.95$
 - $P(A_{1440} \text{ gets me there on time} | \dots) = 0.9999$
- Which action to choose?
- Depends on my preferences for missing flight vs. airport cuisine, etc.
- Utility theory is used to represent and infer preferences
- Decision theory = utility theory + probability theory

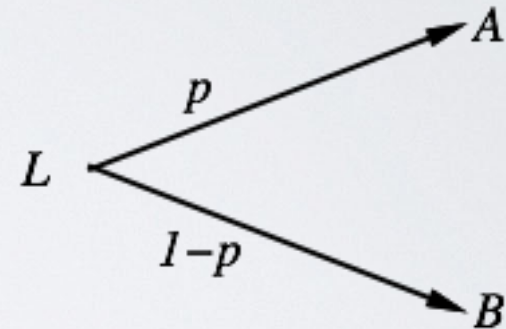
MAKING DECISIONS UNDER UNCERTAINTY

- Rational decision making requires reasoning about one's *uncertainty* and *objectives*
- Previous section focused on uncertainty
- This section will discuss how to make rational decisions based on a *probabilistic model* and *utility function*
- Focus will be on single step decisions, later we will consider sequential decision problems

PREFERENCES

- An agent chooses among prizes (A , B , etc.) and lotteries, i.e., situations with uncertain prizes

- Lottery $L=[p,A; (1-p),B]$



- Notation:

- $A \succ B$ A preferred to B

- $A \sim B$ indifference between A and B

- $A \succeq B$ B not preferred to A

RATIONAL PREFERENCES

- Idea: preferences of a rational agent (*not* a human!) must obey constraints
- Rational preferences \Rightarrow behavior describable as maximization of expected utility
- The Axioms of Rationality:

Orderability

$$(A \succ B) \vee (B \succ A) \vee (A \sim B)$$

Transitivity

$$(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$$

Continuity

$$A \succ B \succ C \Rightarrow \exists p [p, A; 1 - p, C] \sim B$$

Substitutability

$$A \sim B \Rightarrow [p, A; 1 - p, C] \sim [p, B; 1 - p, C]$$

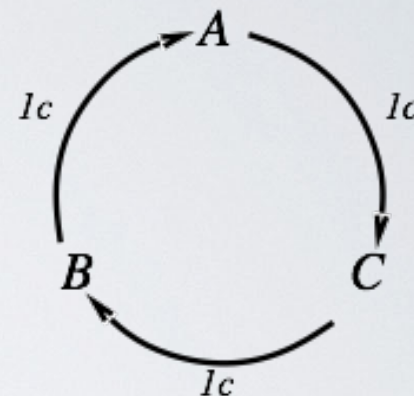
Monotonicity

$$A \succ B \Rightarrow$$

$$(p \geq q \Leftrightarrow [p, A; 1 - p, B] \succeq [q, A; 1 - q, B])$$

RATIONAL PREFERENCES

- Violating the constraints leads to self-evident irrationality
- For example: an agent with intransitive preferences can be induced to give away all its money
- If $B \succ C$, then an agent who has C would pay (say) 1 cent to get B
- If $A \succ B$, then an agent who has B would pay (say) 1 cent to get A
- If $C \succ A$, then an agent who has A would pay (say) 1 cent to get C



REMINDER: EXPECTATION

- The expected value of a function of a random variable is the average, weighted by the probability distribution over outcomes
- Example: How long to get to the airport?

Probability: 0.25 0.50 0.25

Time: 20 min + 30 min + 60 min

35 min

MAXIMIZING EXPECTED UTILITY (MEU)

- Theorem (Ramsey, 1931; von Neumann and Morgenstern, 1944): Given preferences satisfying the constraints there exists a real-valued function U such that
 - $U(A) \geq U(B) \Leftrightarrow A \succeq B$
 - $U(A) > U(B) \Leftrightarrow A \succ B$
 - $U(A) = U(B) \Leftrightarrow A \sim B$
 - $U([p_1, s_1; \dots; p_n, s_n]) = \sum_i p_i U(s_i)$
- *MEU principle*: Choose the action that maximizes expected utility

PREFERENCES LEAD TO UTILITIES

- Note: an agent can be entirely rational (consistent with MEU) without ever representing or manipulating utilities and probabilities
 - E.g., a lookup table for perfect tic-tac-toe
- Although a utility function must exist, it is not unique
 - If $U'(S)=aU(S)+b$ and a and b are constants with $a>0$, then preferences of U' are the same as U
 - E.g., temperatures in Celcius, Fahrenheit, Kelvin

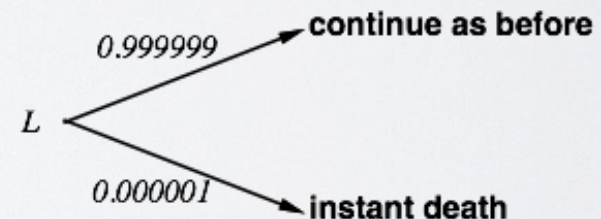
MAXIMIZING EXPECTED UTILITY

- Agent has made some (imperfect) observation o of the state of the world
- If the agent executes action a , the probability the state of the world becomes s' is given by $P(s' | o, a)$
- Preferences on outcomes is encoded using utility function $U(s)$
- Expected utility:
- Principal of maximum expected utility says that a rational agent should choose the action that maximizes expected utility $a^* = \operatorname{argmax}_a EU(a|o)$

UTILITIES: PREFERENCE ELICITATION

- When building a decision-making or decision-support system, it is often helpful to infer the utility function from a human
- Utilities map states to real numbers. Which numbers?
- Standard approach to assessment of human utilities: compare a given state A to a standard lottery L_p that has
 - “best possible prize” u_{\top} with probability p
 - “worst possible catastrophe” u_{\perp} with probability $(1 - p)$
- Adjust lottery probability p until $A \sim L_p$

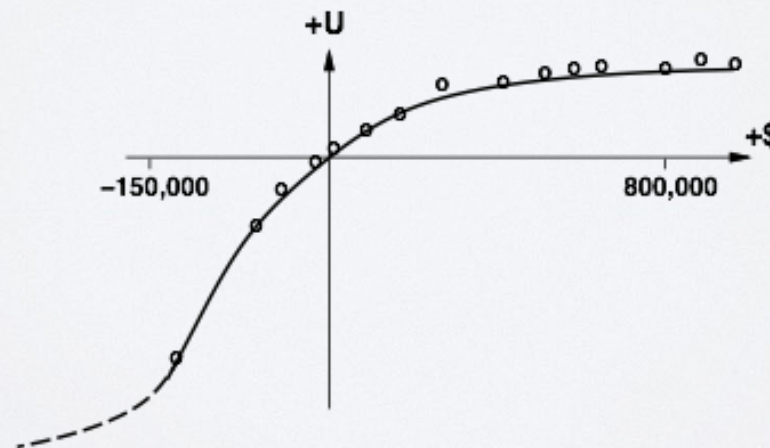
pay \$30 \sim



- Alternatively, set best possible utility to 1 and worst possible to 0

MONEY

- Money does not behave as a utility function
- Given a lottery L with expected monetary value $EMV(L)$, usually $U(L) < U(EMV(L))$, i.e., people are risk-averse
- Utility curve: for what probability p am I indifferent between a prize x and a lottery $[p, \$M; (1-p), \$0]$ for large M ?
- Typical empirical data, extrapolated with risk-prone behavior (utility of money is proportional to the logarithm of the amount):



SUMMARY

- Probability is a rigorous formalism for uncertain knowledge
- Joint probability distribution specifies probability of every atomic event
- Queries can be answered by summing over atomic events
- For nontrivial domains, we must find a way to reduce the joint size
- Independence and conditional independence provide the tools
- Next time: sequential decision making!