

Basic Probability and Decisions

Chris Amato
Northeastern University

Some images and slides are used from: Rob Platt,
CS188 UC Berkeley, AIMA

Uncertainty

Let action A_t = leave for airport t minutes before flight
Will A_t get me there on time?

Problems:

1. partial observability (road state, other drivers' plans, etc.)
2. noisy sensors (traffic reports)
3. uncertainty in action outcomes (flat tire, etc.)
4. immense complexity of modeling and predicting traffic

Hence a purely deterministic (logical) approach either

1) risks falsehood: " A_{25} will get me there on time" or 2) leads to conclusions that are too weak for decision making:

" A_{25} will get me there on time if there's no accident on the bridge and it doesn't rain and my tires remain intact etc etc."

Probability

Probabilistic assertions summarize effects of

- laziness: failure to enumerate exceptions, qualifications, etc.
- ignorance: lack of relevant facts, initial conditions, etc.

Subjective or Bayesian probability:

Probabilities relate propositions to one's own state of knowledge

- e.g., $P(A_{25} | \text{no reported accidents}) = 0.06$

These are not claims of a “probabilistic tendency” in the current situation (but might be learned from past experience of similar situations)


Probabilities of propositions change with new evidence:

- e.g., $P(A_{25} | \text{no reported accidents, 5 a.m.}) = 0.15$

(Discrete) random variables

What is a random variable?

Discrete random variable, X , can take on many (possibly infinite) values, called the *state space* or *domain* $A = \{1, 2, 3, 4, 5, 6\}$ (e.g., a die)



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., $\textit{Temp} < 22.0$


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., $P(X = x) = U[18, 26](x) =$ 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:

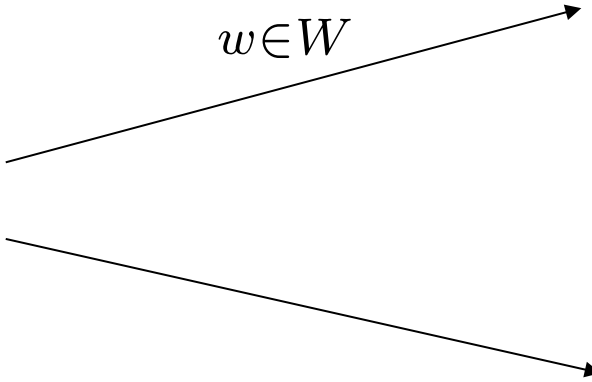
| T | W | $P(T,W)$ |
|------|------|----------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

How many entries do I need here?

Marginalization

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

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

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


| T | P(T) |
|------|------|
| hot | 0.5 |
| cold | 0.5 |

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

| W | P(W) |
|------|------|
| sun | 0.5 |
| rain | 0.4 |

Marginalization

$P(X, Y)$

| X | Y | P |
|----|----|-----|
| +x | +y | 0.2 |
| +x | -y | 0.3 |
| -x | +y | 0.4 |
| -x | -y | 0.1 |



$$P(x) = \sum_y P(x, y)$$



$$P(y) = \sum_x P(x, y)$$

$P(X)$

| X | P |
|----|---|
| +x | |
| -x | |

$P(Y)$

| Y | P |
|----|---|
| +y | |
| -y | |

Conditional Probabilities

Conditional or posterior probabilities

- e.g., $P(\text{cavity}|\text{toothache}) = 0.8$
- i.e., given that toothache is all I know

Notation for conditional distributions: $P(\text{Cavity}|\text{Toothache}) = 2$ -
element vector of 2-element vectors

If we know more, e.g., cavity is also given, then we have $P(\text{cavity}|\text{toothache}, \text{cavity}) = 1$

- Note: the less specific belief remains valid after more evidence arrives, but is not always useful

New evidence may be irrelevant, allowing simplification

- e.g., $P(\text{cavity}|\text{toothache}, \text{redsoxwin}) = P(\text{cavity}|\text{toothache}) = 0.8$

This kind of inference, sanctioned by domain knowledge, is crucial

Conditional Probabilities

Conditional probability: $P(A|B) = \frac{P(A,B)}{P(B)}$ (if $P(B) > 0$)

Example: Medical diagnosis

Product rule: $P(A,B) = P(A \wedge B) = P(A|B)P(B)$

Marginalization with conditional probabilities:

$$P(A) = \sum_{b \in B} P(A|B=b)P(B=b)$$

This formula/rule is called the *law of total probability*

Chain rule is derived by successive application of product rule:

$$\begin{aligned} P(X_1, \dots, X_n) &= P(X_1, \dots, X_{n-1}) P(X_n | X_1, \dots, X_{n-1}) \\ &= P(X_1, \dots, X_{n-2}) P(X_{n-1} | X_1, \dots, X_{n-2}) P(X_n | X_1, \dots, X_{n-1}) = \dots \\ &= \prod_{i=1}^n P(X_i | X_1, \dots, X_{i-1}) \end{aligned}$$

Conditional Probabilities

$P(\textit{sun}|\textit{hot}) \equiv$ Probability that it is sunny *given* that it is hot.


| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

Conditional Probabilities

Calculate the conditional probability using the product rule:

Product rule

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



| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

$$\begin{aligned} P(W = s|T = c) &= \frac{P(W = s, T = c)}{P(T = c)} = \frac{0.2}{0.5} = 0.4 \\ &= P(W = s, T = c) + P(W = r, T = c) \\ &= 0.2 + 0.3 = 0.5 \end{aligned}$$

Conditional Probabilities

$P(X, Y)$

| X | Y | P |
|----|----|-----|
| +x | +y | 0.2 |
| +x | -y | 0.3 |
| -x | +y | 0.4 |
| -x | -y | 0.1 |

- $P(+x \mid +y) ?$

- $P(-x \mid +y) ?$

- $P(-y \mid +x) ?$

Conditional distribution

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

| T | W | $P(T,W)$ |
|------|------|----------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |



| W | $P(W t = hot)$ |
|------|------------------|
| sun | 0.8 |
| rain | 0.2 |

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

Conditional distribution

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

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

| W | P(W $t = hot$) |
|------|-------------------|
| sun | 0.8 |
| rain | 0.2 |

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

$$\begin{aligned} P(sun|hot) &= \frac{P(sun, hot)}{P(hot)} = \frac{P(sun, hot)}{P(sun, hot) + P(rain, hot)} \\ &= \frac{0.4}{0.4 + 0.1} \end{aligned}$$

Conditional distribution

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

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

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

| W | P(W $t = hot$) |
|------|-------------------|
| sun | 0.8 |
| rain | 0.2 |

| W | P(W $t = cold$) |
|------|--------------------|
| sun | 0.4 |
| rain | 0.6 |

Conditional distribution

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

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

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

| W | P(W $t = hot$) |
|------|-------------------|
| sun | 0.8 |
| rain | 0.2 |

| W | P(W $t = cold$) |
|------|--------------------|
| sun | 0.4 |
| rain | 0.6 |

$$\begin{aligned} P(sun|cold) &= \frac{P(sun, cold)}{P(cold)} = \frac{P(sun, cold)}{P(sun, cold) + P(rain, cold)} \\ &= \frac{0.2}{0.2 + 0.3} \end{aligned}$$

Normalization

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

| T | W | P(T,W) |
|---|------|--------|
| h | sun | 0.5 |
| h | rain | 0.3 |
| c | sun | 0.2 |
| c | rain | 0.0 |

| W | P(W $t = hot$) |
|------|-------------------|
| sun | 0.8 |
| rain | 0.2 |

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


| W | P(W $t = cold$) |
|------|--------------------|
| sun | 0.4 |
| rain | 0.6 |

Can we avoid explicitly computing this?


$$\begin{aligned}
 P(sun|cold) &= \frac{P(sun, cold)}{P(cold)} = \frac{P(sun, cold)}{P(sun, cold) + P(rain, cold)} \\
 &= \frac{0.2}{0.2 + 0.3}
 \end{aligned}$$

Normalization

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |



| W | P(W, t=hot) |
|------|-------------|
| sun | 0.4 |
| rain | 0.1 |



| W | P(W t = hot) |
|------|----------------|
| sun | 0.8 |
| rain | 0.2 |


Select corresponding elements
from the joint distribution

Scale the numbers so
that they sum to 1


$$P(\text{sun}|\text{cold}) = \frac{P(\text{sun}, \text{cold})}{P(\text{cold})} = \frac{P(\text{sun}, \text{cold})}{P(\text{sun}, \text{cold}) + P(\text{rain}, \text{cold})}$$

Normalization

| T | W | P(T,W) |
|------|------|--------|
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |



| W | P(W, t=hot) |
|------|-------------|
| sun | 0.4 |
| rain | 0.1 |



| W | P(W t = hot) |
|------|----------------|
| sun | 0.8 |
| rain | 0.2 |

Select corresponding elts from the joint distribution

Scale the numbers so that they sum to 1.

$$P(\text{sun}|\text{cold}) = \frac{P(\text{sun}, \text{cold})}{P(\text{cold})} = \frac{P(\text{sun}, \text{cold})}{P(\text{sun}, \text{cold}) + P(\text{rain}, \text{cold})}$$

The only purpose of this denominator is to make the distribution sum to one.


– we achieve the same thing by scaling.

Normalization


$P(X \mid Y=-y) ?$

$P(X, Y)$

| X | Y | P |
|----|----|-----|
| +x | +y | 0.2 |
| +x | -y | 0.3 |
| -x | +y | 0.4 |
| -x | -y | 0.1 |



?



?


Independence

A and B are *independent* iff

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

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

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



32 entries reduced to 12; for n independent biased coins, $2^n \rightarrow n$

Absolute independence powerful but rare

Dentistry is a large field with hundreds of variables, none of which are independent. What to do?

Conditional independence

$P(\text{Toothache}, \text{Cavity}, \text{Catch})$ has $2^3 - 1 = 7$ independent entries

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

$$(1) P(\text{catch}|\text{toothache}, \text{cavity}) = P(\text{catch}|\text{cavity})$$

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

$$(2) P(\text{catch}|\text{toothache}, \neg\text{cavity}) = P(\text{catch}|\neg\text{cavity})$$

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 (equations 1 and 2 remove 2)

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?

Bayes' Rule

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

It's easy to derive from the product rule:

$$P(a, b) = P(b|a)P(a) = \underbrace{P(a|b)}_{\substack{\uparrow \\ \text{Solve for this}}} P(b)$$

Solve for this

Using Bayes' Rule

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



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

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

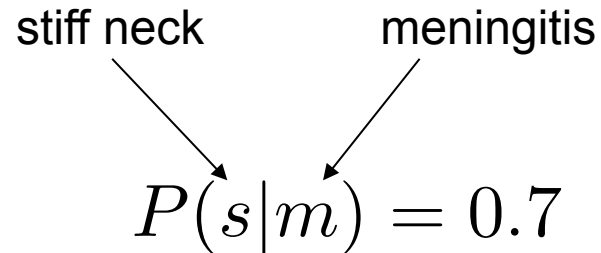
It's often easier to estimate this

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

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

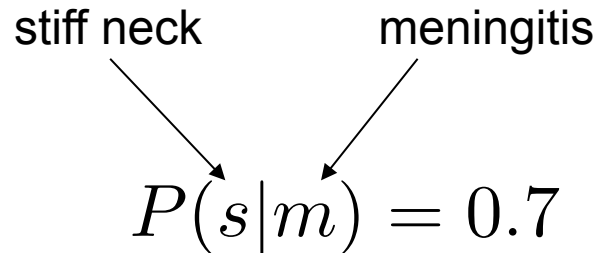
What are the chances that you have meningitis?

Bayes' Rule Example

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

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

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

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

$$P(s) = 0.01$$



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:

$P(W)$

| R | P |
|------|-----|
| sun | 0.8 |
| rain | 0.2 |

$P(D|W)$

| D | W | P |
|-----|------|-----|
| wet | sun | 0.1 |
| dry | sun | 0.9 |
| wet | rain | 0.7 |
| dry | rain | 0.3 |


- What is $P(W|dry)$?

Bayes' rule and conditional independence

$$\begin{aligned} P(\text{Cavity}|\text{toothache},\text{catch}) \\ &= \alpha P(\text{toothache},\text{catch}|\text{Cavity})P(\text{Cavity}) \\ &= \alpha P(\text{toothache}|\text{Cavity})P(\text{catch}|\text{Cavity})P(\text{Cavity}) \end{aligned}$$

This is an example of a naive Bayes model:

$$P(\text{Cause}, \text{Effect}_1, \dots, \text{Effect}_n) = P(\text{Cause})\prod_i P(\text{Effect}_i|\text{Cause})$$



Total number of parameters is linear in n

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?

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, next week 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 $B > C$, then an agent who has C would pay (say) 1 cent to get B

If $A > B$, then an agent who has B would pay (say) 1 cent to get A

If $C > A$, then an agent who has A would pay (say) 1 cent to get C






Reminder: Expectations

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?

| | | | | | | | |
|--------------|--------|---|--------|---|--------|--|--------|
| Time: | 20 min | | 30 min | | 60 min | | |
| | x | + | x | + | x | | |
| Probability: | 0.25 | | 0.50 | | 0.25 | | 35 min |



Mean and variance

Mean, μ , or expected value:

Discrete:
$$\mathbb{E}[X] = \sum_{x \in X} xP(x)$$

Continuous:
$$\mathbb{E}[X] = \int_x xP(x)dx$$

Variance:
$$\text{var}[X] = \mathbb{E}[(X - \mu)^2]$$

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

MEU continued

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: $EU(a | o) = \sum_{s'} P(s' | a, o) U(s')$

Principle 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 ~



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

Utility scales

Normalized utilities: $u_{\top} = 1.0$, $u_{\perp} = 0.0$

Micromorts: one-millionth chance of death

Useful for Russian roulette, paying to reduce product risks, etc.

QALYs: quality-adjusted life years

Useful for medical decisions involving substantial risk

Note: behavior is invariant w.r.t. positive linear transformation

$$U'(x) = k_1 U(x) + k_2 \quad \text{where } k_1 > 0$$

With deterministic prizes only (no lottery choices), only ordinal utility can be determined, i.e., total order on prizes


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):



Student group utility

Who prefers the lottery at different values of p ? ($M=10,000$)



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!