Artificial Intelligence

Uncertainty & probability
Chapter 13, AIMA

”When an agent knows enough facts about its environment, the logical approach enables it to derive plans that are guaranteed to work.”

”Unfortunately, agents almost never have access to the whole truth about the environment.”

From AIMA, chapter 13

Airport example

Let action A_t = leave for airport *t* minutes before the flight.

Will a given A_t get me there on time?

Problems:

- partial observability (road state, other drivers' plans, car condition, ...)
- noisy sensors (maps, KCBS traffic reports, news, ...)
- uncertainty in action outcomes (flat tire, missing an exit, ...)
- immense complexity in modelling and predicting state of the environment

Hence a purely logical approach either:

1) risks falsehood: *"A_{50} will get me there on time"*, or
2) leads to conclusions that are too weak for decision making:

"A_{50} will get me there on time if there's no accident and it doesn't rain and my tires remain intact and ..."

(A_{1440} might reasonably be said to get me to there on time but I'd have to stay overnight in the airport... it's not a good solution, even if it is "correct" in some sense)

Slide adapted from S. Russell

Why FOL fails?

Laziness: People who want to solve problems can’t be bothered to list all possible requirements

- FOL is often the very definition of „useless solution”

Ignorance: We have no theoretical models for the domain that would be complete

- and what we have isn't usually tested well enough

Practical reasons: Even if we knew everything and could be persuaded to list everything, the rules would be totally impractical to use

- we can’t possibly account for / sense / test everything

Instead, use decision theory

Decision theory = probability theory + utility theory

Probability

- Assign probability to a proposition based on the percepts, i.e. the information the agent has.
- Proposition is either true or false. Probability means assigning a value that indicates how much do we believe in it being true/false.
- *Evidence* is all information that the agent receives. Probabilities can (will) change when more evidence is acquired.
- Prior/unconditional probability $\Leftrightarrow$ no evidence at all.
- Posterior/conditional probability $\Leftrightarrow$ after evidence is obtained.

Utility

- Plan does not need to be guaranteed to achieve the goal.
- To make choices, the agent must have preferences between the different possible outcomes of various plans.
- Utility represents the value of the outcomes, and utility theory is used to reason with the resulting preferences.
- An agent is rational if and only if it chooses the action that yields the highest expected utility, when averaged over all possible outcomes.

Inspired by Sun-Hwa Hahn

Probability

The dealer will pay $1 if you flip a coin and it lands head up...

How much will you pay to play this game?

Most people agree that they would pay up to $0.50

Probability

The dealer will pay $2 if you roll a die and it lands with a 6 up...

How much will you pay to play this game?

Most people agree that they would pay up to $0.33

Probability

The dealer will pay \$2 if the card you draw has a rank at least as high as the rank of the card he draws....

How much will you pay to play this game?

$$52*3 + (52*51 - 52*3) / 2 = 52 * 27$$

$$(52 * 27) / (52 * 51) = 27 / 51$$

$$\$2 * 27 / 51 = \$1.0588$$

Symmetry

In principle, according to classical physics, we should be able to predict how the coin will land, if we know its initial position, the force exerted upon it by the flipper, the position of the surface on which it lands, the material properties of the coin and that surface, the air pressure, any winds that blow through the trajectory of the coin, etc.

We can apply Newton's laws of motion and the law of gravity, while accounting for the elasticity of the collisions between the coin and the surface, and for air friction, and for whatever other physical effects there are, to calculate the coin's motion until it stops.

In practice, of course, these calculations are too difficult for us to do exactly, at least in our heads, even if we knew all the relevant factors.

And furthermore, we do not know all of these factors.

Slide from David A. Meyer

Basic Principle of Counting

But there is a symmetry in the problem:
none of these factors differ significantly with the different sides of the coin

Thus the amount we are willing to pay to win $1 if the coin lands *head up* is the same amount that we are willing to pay to win $1 if the coin lands *tail up*

We believe that coin will land head up or tail up
$p + p = 1, \; p = 0.5$
We believe that a die will show a number 1..6
$6 * p = 1, \; p = 0.1666666666(6)$

Interpretation of Probability

Classical interpretation

The probability of an event is the ratio of the number of cases favorable to it, to the number of all cases possible when nothing leads us to expect that any one of these cases should occur more than any other

Frequency interpretation

The probability of an event is the limit of its relative frequency in an infinitely large number of trials

Bayesian interpretation

The probability of an event depends on some prior probability, which is then updated in the light of new relevant data or observations

Boolean variable

$$X \in \{True, False\}$$

$$P(X = true) + P(X = false) = 1$$

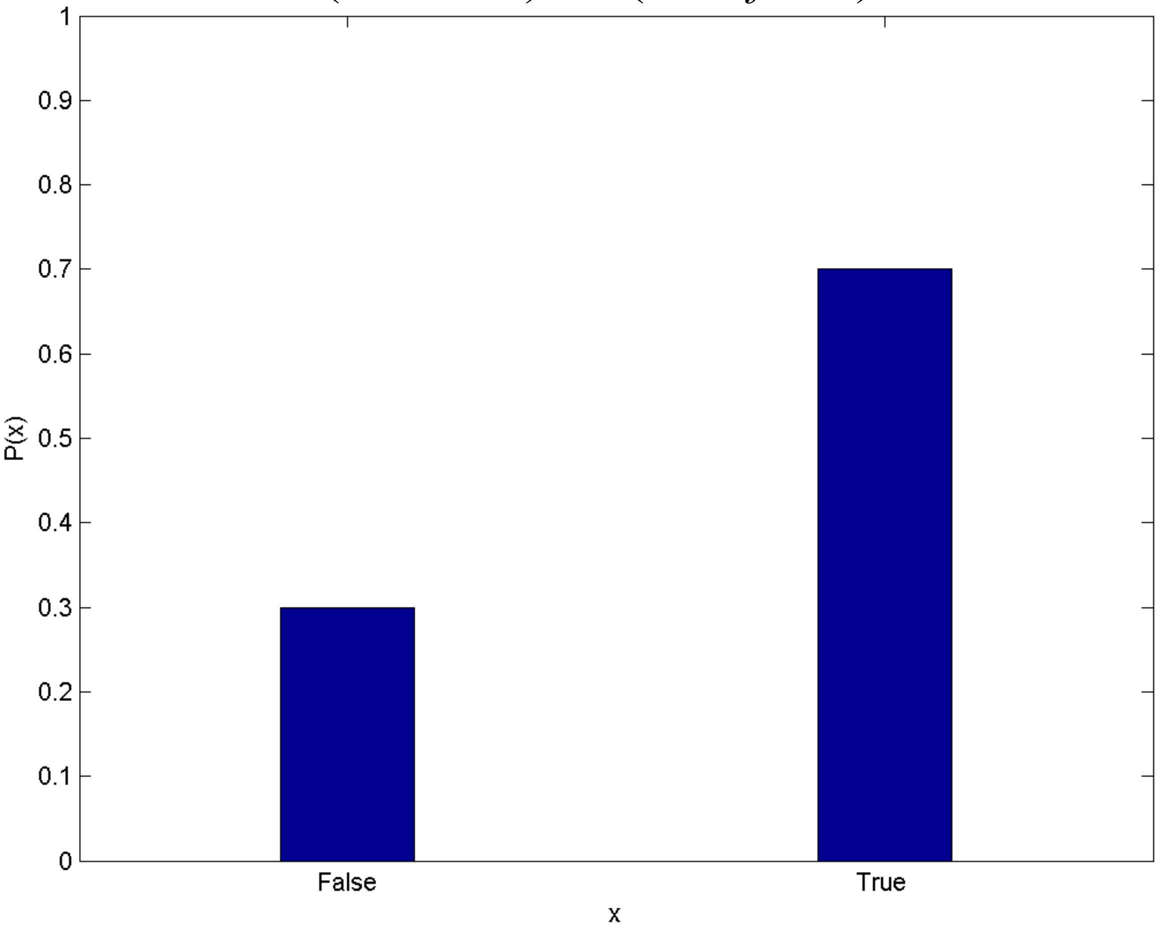

$P(X)=0.7$
$P(\neg X)=0.3$

Examples:

W_{31}
Smell
CatchFlight
Cancer
Cavity
LectureToday
...

Discrete variable

$$X \in \{x_1, x_2, x_3, x_4, \ldots\}$$

$$\sum_i P(X = x_i) = 1$$

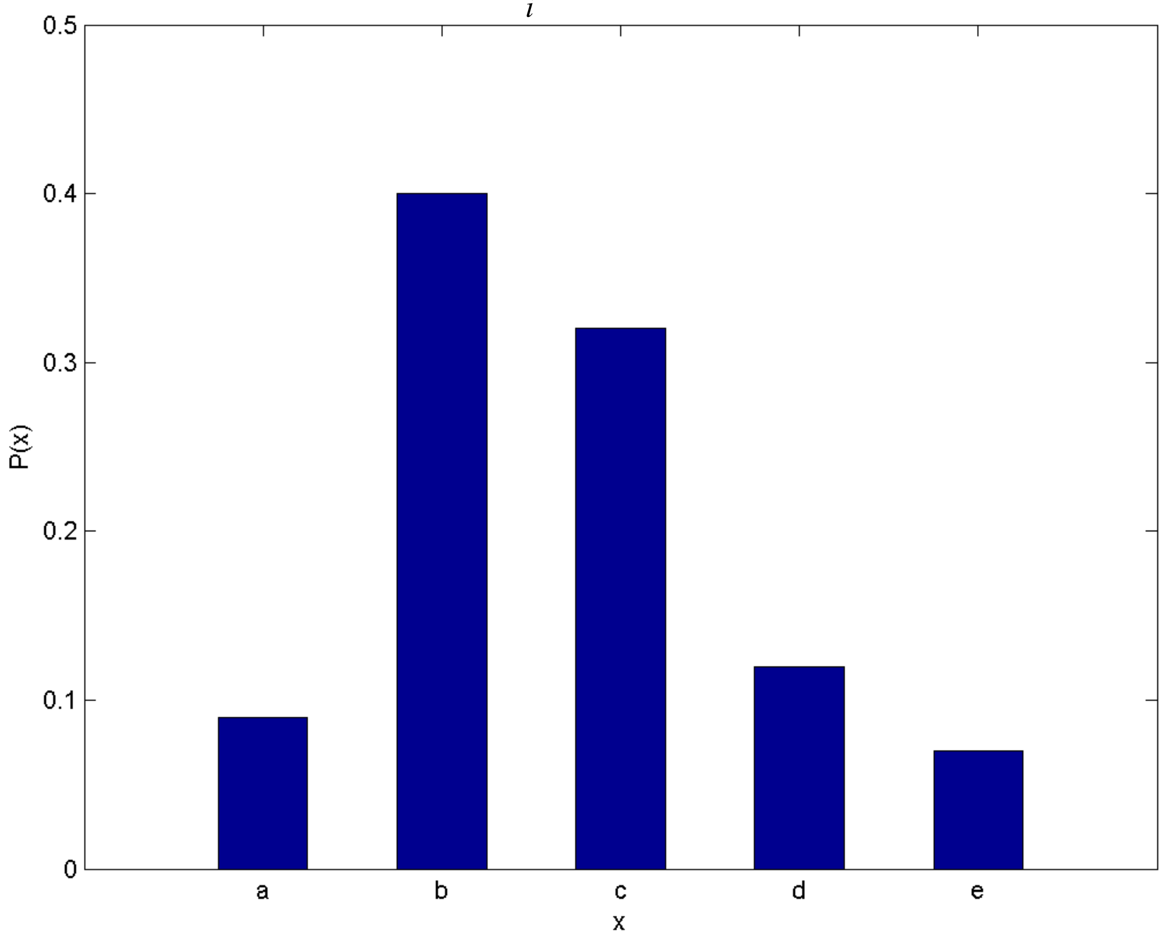

$P(a)=0.1$
$P(b)=0.4$
$P(c)=0.3$
$P(d)=0.12$
$P(e)=0.07$

Examples:

Wumpus
Weather
Color
Age
LectureQuality
PlayingCard
...

Continuous variable X

Probability density $P(x)$

$$\int_{-\infty}^{\infty} P(x)dx = 1$$

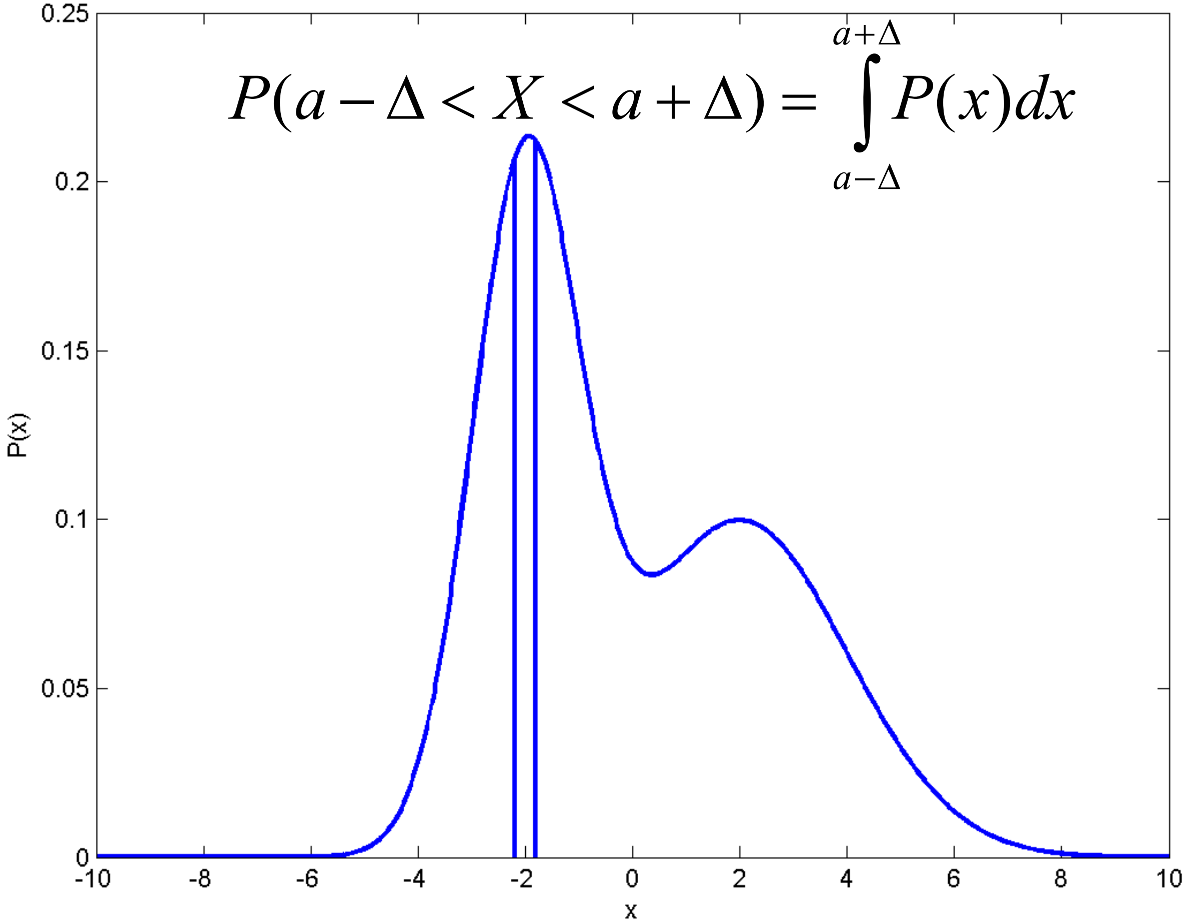

Examples:
Temperature
WindSpeed
OxygenLevel
LectureLength
...

Probability Axioms

The probability of an event is a non-negative real number:

$$P(E) \in \mathbb{R} \wedge P(E) \geq 0 \qquad \forall E \in F$$

The probability that *some* elementary event in the entire sample space will occur is 1:

$$P(\Omega) = 1$$

Any countable sequence of pairwise disjoint (i.e. mutually exclusive) events $E_1, E_2, \ldots$ satisfies:

$$P(E_1 \cup E_2 \cup \cdots) = \sum_{i=1}^{\infty} P(E_i).$$

From Wikipedia

Properties of Probability

Monotonicity:

$$P(A) \leq P(B) \quad \text{if} \quad A \subseteq B.$$

The probability of the empty set:

$$P(\emptyset) = 0.$$

The numeric bound:

$$0 \leq P(E) \leq 1 \qquad \text{for all } E \in F.$$

The addition law of probability:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

The inclusion-exclusion principle:

$$P(\Omega \setminus E) = 1 - P(E)$$

From Wikipedia

Expected Value

$$\mathrm{E}[X] = x_1 p_1 + x_2 p_2 + \ldots + x_k p_k \; .$$

$$\mathrm{E}[X] = \sum_{i=1}^{\infty} x_i \, p_i,$$

$$\mathrm{E}[X] = \int_{-\infty}^{\infty} x f(x) \, \mathrm{d}x.$$

From Wikipedia

Propositions

Elementary:

Cavity = True, W_{31} = True, Cancer = False, Card = A , Card = 4 , Weather = Sunny, Age = 40, (20°C < Temperature < 21°C), (2 hrs < LengthLecture < 3 hrs)

Complex: (Elementary + connective)

$\neg$Cancer $\wedge$ Cavity

$Card_1$ = A $\wedge$ $Card_2$ = A

Sunny $\wedge$ (30°C < Temperature) $\wedge$ $\neg$Beer

Propositions

Cavity = True, W_{31} = True, Cancer = False
2 hrs < LengthLecture < 3 hrs
$Card_1 = A \quad \wedge Card_2 = A$
Sunny $\wedge$ (30°C < Temperature) $\wedge$ $\neg$Beer

In the physical world, every of those propositions is either true or false

– we are not talking about degrees of truth here (that's called *fuzzy logic)*

But we can discuss different degrees to which we believe that various propositions are true or not

Extension of propositional calculus...

Atomic event

is a complete specification of one of the possible states of the world/environment

- Mutually exclusive
- Exhaustive

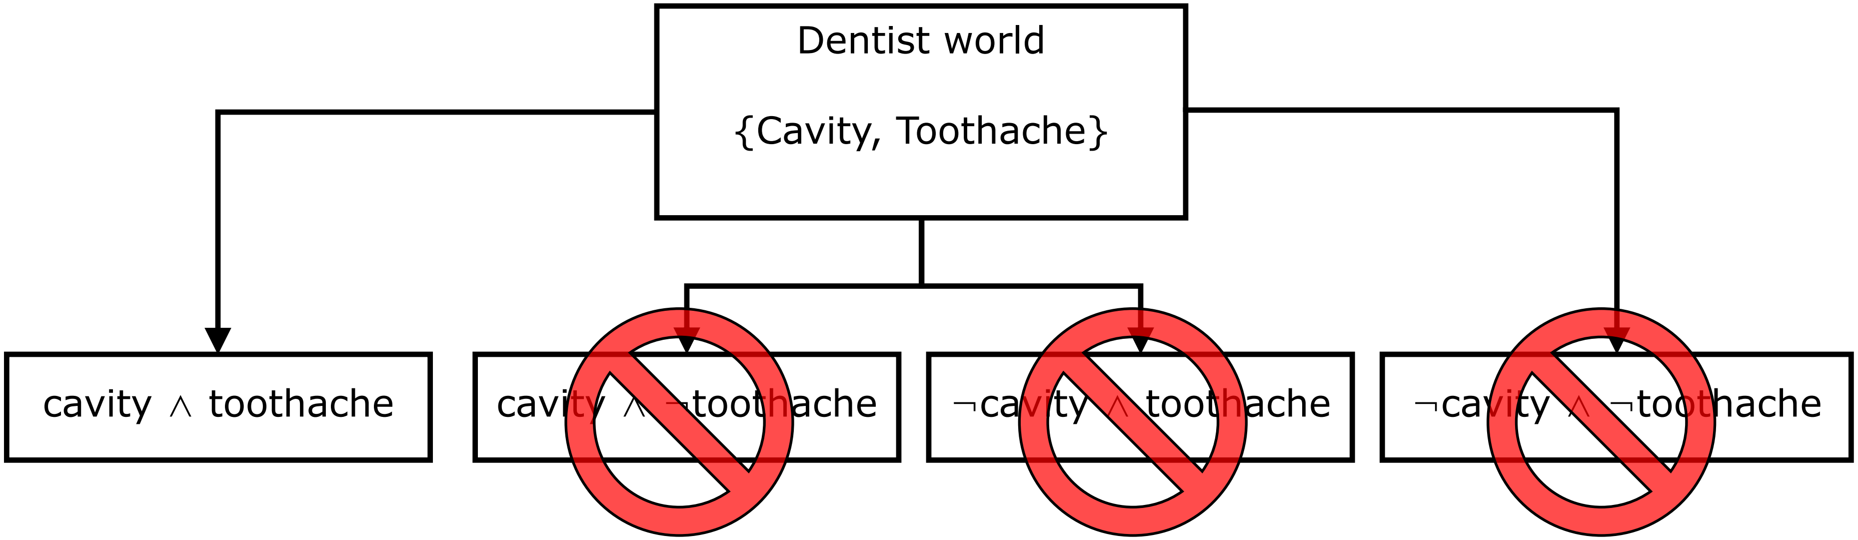

Prior & posterior

- **Prior probability** $P(X=a) \equiv P(a)$
  Our belief in X=a being true before any information is collected

- **Posterior probability** $P(X=a|Y=b) \equiv P(a|b)$
  Our belief that X=a is true when we know that Y=b is true (and we don't know anything else)

- **Joint probability** $P(X=a, Y=b) \equiv P(a \wedge b) \equiv P(a,b)$
  Our belief that $(X=a \wedge Y=b)$ is true.

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

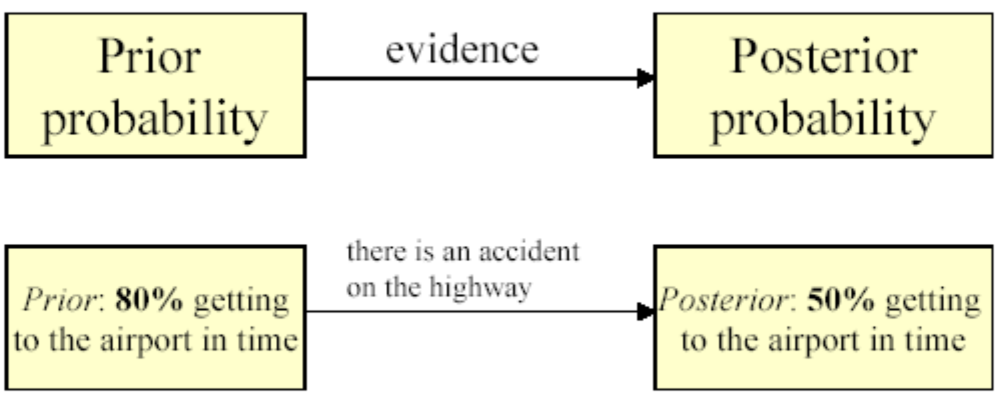

Image borrowed from Lazaric & Sceffer

Conditional probability
(Venn diagram)

$$P(a|b) = P(a,b) / P(b)$$

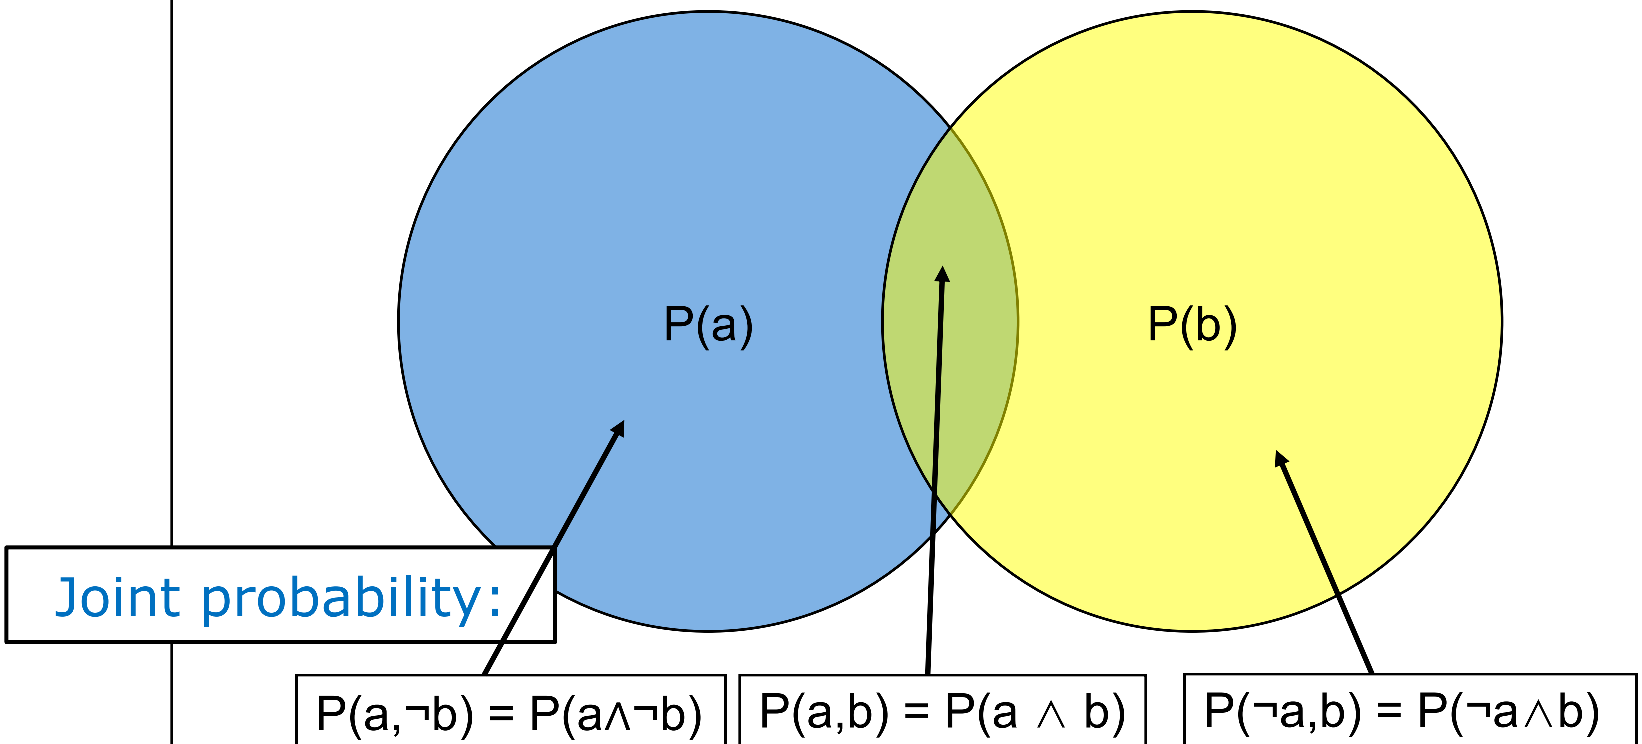

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the conditional probability (risk) for failing the course if the student fails the first two attempts at the oral exam?

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

$$P(2 = \text{ace} \mid 1 = \text{ace}) = \frac{3}{51}$$

What's the probability that both are aces?

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the conditional probability (risk) for failing the course if the student fails the first two attempts at the oral exam?

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

$$P(2 = ace|1 = ace) = \frac{3}{51} \qquad P(2 = \text{ace}, 1 = \text{ace}) = P(2 = \text{ace} \mid 1 = \text{ace})P(1 = \text{ace}) = \frac{3}{51} \cdot \frac{4}{52} = 0.5\%$$

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the conditional probability (risk) for failing the course if the student fails the first two attempts at the oral exam?

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

$$P(2=\text{ace}\,|\,1=\text{ace})=\frac{3}{51} \qquad P(2=\text{ace},1=\text{ace})=P(2=\text{ace}\,|\,1=\text{ace})P(1=\text{ace})=\frac{3}{51}\cdot\frac{4}{52}=0.5\%$$

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the conditional probability (risk) for failing the course if the student fails the first two attempts at the oral exam?

$$P(\text{Fail course}\,|\,\text{Fail OE 1 \& 2})=\frac{P(\text{Fail course, Fail OE1 \& 2})}{P(\text{Fail OE1 \& 2})}=\frac{8/63}{15/63}=53\%$$

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

$$63=23+25+7+8$$

$$15=63-23-25$$

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

$$P(2=\text{ace}\,|\,1=\text{ace}) = \frac{3}{51} \qquad P(2=\text{ace},1=\text{ace}) = P(2=\text{ace}\,|\,1=\text{ace})P(1=\text{ace}) = \frac{3}{51}\cdot\frac{4}{52} = 0.5\%$$

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the cor[...]he student fails the firs[...]

What's the prior probability for not passing the course?

$$P(\text{Fail course}\,|\,\text{Fail OE 1 \& 2}) = \frac{P(\text{Fail course, Fail OE1 \& 2})}{P(\text{Fail OE1 \& 2})} = \frac{8/63}{15/63} = 53\%$$

$\boxed{8/63 \ = 13\%}$

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

Conditional probability examples

1. You draw two cards randomly from a deck of 52 playing cards. What is the conditional probability that the second card is an ace if the first card is an ace?

$$P(2=\text{ace}\mid 1=\text{ace})=\frac{3}{51} \qquad P(2=\text{ace},1=\text{ace})=P(2=\text{ace}\mid 1=\text{ace})P(1=\text{ace})=\frac{3}{51}\cdot\frac{4}{52}=0.5\%$$

1. In a course I'm giving with oral examination the examination statistics over the period 2002-2005 have been: 23 have passed the oral examination in their first attempt, 25 have passed it their second attempt, 7 have passed it in their third attempt, and 8 have not passed it at all (after having failed the oral exam at least three times). What is the conditional probability (risk) for failing the course if the student fails the first two attempts at the oral exam?

$$P(\text{Fail course}\mid\text{Fail OE 1 \& 2})=\frac{P(\text{Fail course, Fail OE1 \& 2})}{P(\text{Fail OE1 \& 2})}=\frac{8/63}{15/63}=53\%$$

1. (2005) 84% of the Swedish households have computers at home, 81% of the Swedish households have both a computer and internet. What is the probability that a Swedish household has internet if we know that it has a computer?

$$P(\text{Internet}\mid\text{Computer})=\frac{P(\text{Internet, Computer})}{P(\text{Computer})}=\frac{0.81}{0.84}=96\%$$

Inference

- Inference means computing

  P(State of the world | Observed evidence)

  P(Y | **e**)

For example: The probability for having a cavity if I have a toothache

$$P(cavity \mid toothache)$$

Or

The probability for having a cavity if I have a toothache and the dentist finds a catch in my tooth during inspection

$$P(cavity \mid toothache, catch)$$

Inference with full joint distribution

- The **full joint probability distribution** is the probability distribution for all random variables used to describe the world.

Dentist example {Toothache, Cavity, Catch}

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Inference with full joint distribution

- The **full joint probability distribution** is the probability distribution for all random variables used to describe the world.

Dentist example {Toothache, Cavity, Catch}

	toothache		¬toothache		
	catch	¬catch	catch	¬catch	
cavity	0.108	0.012	0.072	0.008	0.200
¬cavity	0.016	0.064	0.144	0.576	

$P(cavity) = 0.2$

Marginal probability (marginalization)

Inference with full joint distribution

- The **full joint probability distribution** is the probability distribution for all random variables used to describe the world.

Dentist example {Toothache, Cavity, Catch}

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576
	0.200			

P(cavity) = 0.2

Marginal probability
(marginalization)

P(toothache) = 0.2

Inference with full joint distribution

- The **full joint probability distribution** is the probability distribution for all random variables used to describe the world.

Dentist example {Toothache, Cavity, Catch}

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

$$\mathbf{P}(\mathbf{Y}) = \sum_{\mathbf{z}} \mathbf{P}(\mathbf{Y}, \mathbf{z})$$

$$\mathbf{P}(\mathbf{Y}) = \sum_{\mathbf{z}} \mathbf{P}(\mathbf{Y} \mid \mathbf{z})\mathbf{P}(\mathbf{z})$$

Marginal probability (marginalization)

$$P(cavity \mid toothache) = \frac{P(cavity \wedge toothache)}{P(toothache)} = \frac{\boxed{0.12}}{\boxed{0.2}} = 0.6$$

$$P(\neg cavity \mid toothache) = \frac{P(\neg cavity \wedge toothache)}{P(toothache)} = \frac{0.08}{0.2} = 0.4$$

$$P(cavity \mid \neg toothache) = \frac{P(cavity \wedge \neg toothache)}{P(\neg toothache)} = \frac{0.08}{0.8} = 0.1$$

Dentist example {Toothache, Cavity, Catch}

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Inference with full joint distribution

The general inference procedure

$$\mathbf{P}(Y \mid \mathbf{e}) = \alpha \mathbf{P}(Y, \mathbf{e}) = \alpha \sum_{z} \mathbf{P}(Y, \mathbf{e}, z)$$

where α is a normalization constant. This can always be computed if the full joint probability distribution is available.

$$\mathbf{P}(Cavity \mid toothache) = \alpha\left[\mathbf{P}(Cavity \mid toothache, catch) + \mathbf{P}(Cavity \mid toothache, \neg catch)\right]$$

Completely impossible to do in practice: $O(2^n)$ complexity

Independence

Independence between variables can dramatically reduce the amount of computation.

$$\mathbf{P}(X,Y) = \mathbf{P}(X)\mathbf{P}(Y)$$

$$\mathbf{P}(X \mid Y) = \mathbf{P}(X)$$

$$\mathbf{P}(Y \mid X) = \mathbf{P}(Y)$$

We don’t need to mix independent variables in our computations

Independence for dentist example

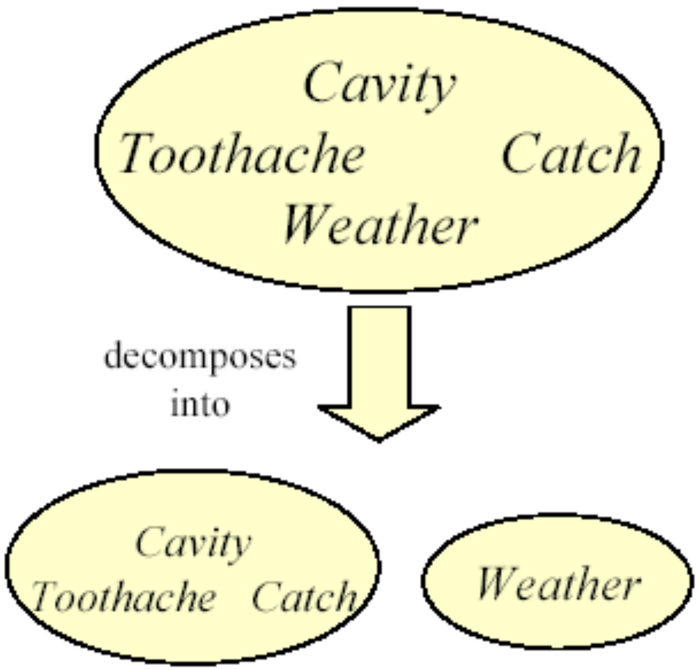

Oral health is independent of weather

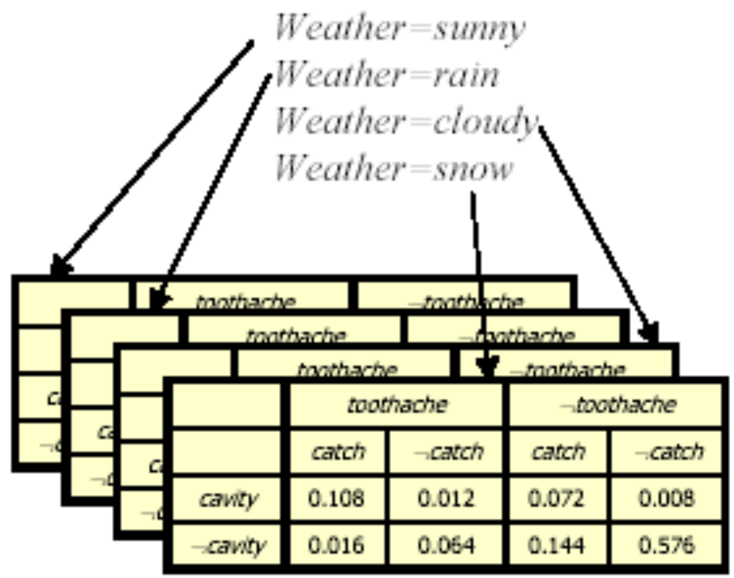

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

decomposes
into

	Weather
sunny	0.4
rain	0.15
cloudy	0.4
snow	0.05

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Image borrowed from Lazaric & Sceffer

Bayes' theorem

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

$$\boxed{P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}}$$

Bayes theorem example

Joe is a randomly chosen member of a large population in which 3% are heroin users. Joe tests positive for heroin in a drug test that correctly identifies users 95% of the time and correctly identifies nonusers 90% of the time.

Is Joe a heroin addict?

Example from http://plato.stanford.edu/entries/bayes-theorem/supplement.html

Bayes theorem example

Joe is a randomly chosen member of a large population in which 3% are heroin users. Joe tests positive for heroin in a drug test that correctly identifies users 95% of the time and correctly identifies nonusers 90% of the time.

Is Joe a heroin addict?

$$P(H \mid pos) = \frac{P(pos \mid H)P(H)}{P(pos)}$$

$$P(H) = 3\% = 0.03,\ P(\neg H) = 1 - P(H) = 0.97$$

$$P(pos \mid H) = 95\% = 0.95,\ P(pos \mid \neg H) = 10\% = 1 - 0.90$$

$$P(pos) = P(pos \mid H)P(H) + P(pos \mid \neg H)P(\neg H) = 0.1255$$

$$P(H \mid pos) = 0.227 \approx 23\%$$

Example from http://plato.stanford.edu/entries/bayes-theorem/supplement.html

Bayes theorem: The Monty Hall Game show

In a TV Game show, a contestant selects one of three doors; behind one of the doors there is a prize, and behind the other two there are no prizes.

After the contestant select a door, the game-show host opens one of the remaining doors, and reveals that there is no prize behind it. The host then asks the contestant whether he/she wants to SWITCH to the other unopened door, or STICK to the original choice.

What should the contestant do?

See http://www.io.com/~kmellis/monty.html

The Monty Hall Game Show

$$\text{prize behind door} = \{1,2,3\}, \;\; \text{open}_i = \text{Host opens door } i$$

The Monty Hall Game Show

prize behind door $= \{1,2,3\}$, open_i = Host opens door i

Contestant selects door 1

Host opens door 2 $\Rightarrow \text{open}_2$

$$P(1 \mid \text{open}_2) = \frac{P(\text{open}_2 \mid 1) P(1)}{P(\text{open}_2)} =$$

$$P(3 \mid \text{open}_2) = \frac{P(\text{open}_2 \mid 3) P(3)}{P(\text{open}_2)} =$$

© Let's make a deal (A Joint Venture)

The Monty Hall Game Show

prize behind door $= \{1,2,3\}$, open_i = Host opens door i

Contestant selects door 1

Host opens door 2 $\Rightarrow \text{open}_2$

© Let's make a deal (A Joint Venture)

$$P(1 \mid \text{open}_2) = \frac{P(\text{open}_2 \mid 1)P(1)}{P(\text{open}_2)} = 1/3$$

$$P(3 \mid \text{open}_2) = \frac{P(\text{open}_2 \mid 3)P(3)}{P(\text{open}_2)} = 2/3$$

$$P(\text{open}_2) = \sum_{i=1}^{3} P(\text{open}_2 \mid i)P(i) = 1/2$$

$$P(\text{open}_2 \mid 1) = 1/2,\ P(\text{open}_2 \mid 2) = 0,\ P(\text{open}_2 \mid 3) = 1$$

Bayes theorem: The Monty Hall Game show

In a TV Game show, a contestant selects one of three doors; behind one of the doors there is a prize, and behind the other two there are no prizes.

After the contestant select a door, the game-show host opens one of the remaining doors, and reveals that there is no prize behind it. The host then asks the contestant whether he/she wants to SWITCH to the other unopened door, or STICK to the original choice.

What should the contestant do?

The host is actually asking the contestant whether he/she wants to SWITCH the choice to both other doors, or STICK to the original choice. Phrased this way, it is obvious what the optimal thing to do is.

Conditional independence

We say that X and Y are conditionally independent given Z iff

$$\mathbf{P}(X,Y \mid Z) = \mathbf{P}(X \mid Z)\mathbf{P}(Y \mid Z)$$

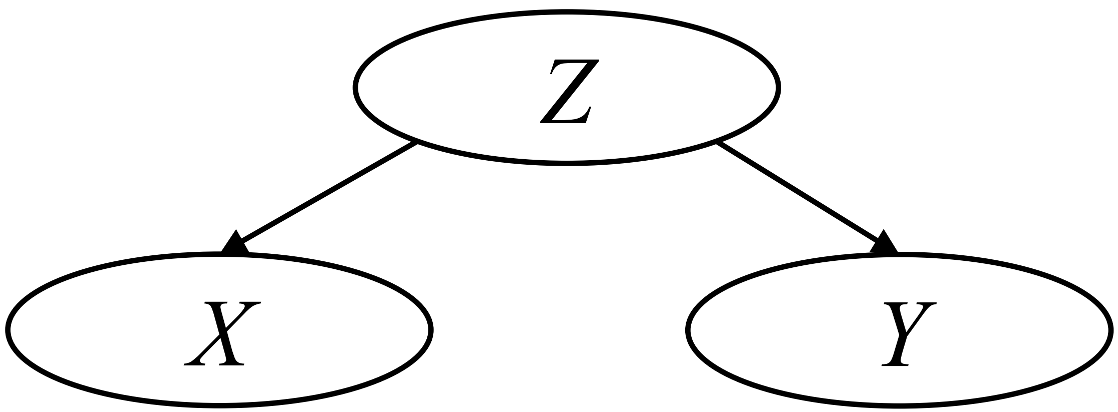

What's the relation between independence and conditional independence?

Naive Bayes: Combining evidence

Assume full conditional independence and express the full joint probability distribution as:

$$\mathbf{P}(Effect_1, Effect_2, \ldots, Effect_n, Cause) =$$

$$\mathbf{P}(Effect_1, Effect_2, \ldots, Effect_n \mid Cause)\mathbf{P}(Cause) =$$

$$\mathbf{P}(Effect_1 \mid Cause) \cdots \mathbf{P}(Effect_n \mid Cause)\mathbf{P}(cause) =$$

$$\left[\prod_{i=1}^{n} \mathbf{P}(Effect_i \mid Cause)\right]\mathbf{P}(Cause)$$

Naive Bayes: Dentist example

$$\mathbf{P}(Toothache, Catch, Cavity) =$$
$$\mathbf{P}(Toothache, Catch \mid Cavity)\mathbf{P}(Cavity) \approx$$
$$\mathbf{P}(Toothache \mid Cavity)\mathbf{P}(Catch \mid Cavity)\mathbf{P}(Cavity)$$
$$\Rightarrow \mathbf{P}(toothache, catch, cavity) \approx$$
$$\frac{(0.108 + 0.012)}{0.2} \times \frac{(0.108 + 0.072)}{0.2} \times 0.2 = 0.108$$
$$\text{True value}: \mathbf{P}(toothache, catch, cavity) = 0.108$$

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Naive Bayes: Dentist example

P(Toothache|Cavity)	toothache	¬toothache
cavity	0.6	0.4
¬cavity	0.1	0.9

2 independent numbers

P(Catch|Cavity)	catch	¬catch
cavity	0.9	0.1
¬cavity	0.2	0.8

2 independent numbers

	P(Cavity)
cavity	0.2
¬cavity	0.8

1 independent number

P(Catch, Toothache,Cavity)	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Full table has 2^3-1=7 independent numbers [$O(2^n)$]

Conclusions

- Uncertainty in knowledge arises because of laziness and ignorance
  - this is unavoidable in the real world
- Probability expresses the degree to which the agent believes in possibility of different outcomes
- Decision theory helps the agent to act in a rational manner in the face of uncertainty
  - rational = maximize it's own expected utility
- In order to get efficient algorithms to deal with probability we need to explore the concept of conditional independence among variables

Artificial Intelligence

Utility theory

Chapter 16, AIMA

The utility function $U(S)$

- An agent's preferences between different states S in the world are captured by the *utility function* $U(S)$.

- If $U(S_i) > U(S_j)$ then the agent prefers state S_i over state S_j
- If $U(S_i) = U(S_j)$ then the agent is indifferent between the two states S_i and S_j

Maximize expected utility

A rational agent should choose the action that maximizes the the agent's expected utility (EU):

$$EU(A \mid \mathbf{E}) = \sum_i P(Result_i(A) \mid Do(A), \mathbf{E}) U(Result_i(A))$$

Where $Result_i(A)$ enumerates all the possible resulting states after doing action A given observed environment state (evidence) E.

The basis of utility theory

Notation:

$A \succ B$ A is preferred to B

$A \sim B$ The agent is indifferent between A and B

$A \succsim B$ The agent prefers A to B, or is indifferent between them.

A lottery is described with

$$L = [p_1, C_1; p_2, C_2; \ldots; p_n, C_n]$$

The six axioms of utility theory

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	$(A \succ B \succ C) \Rightarrow \exists p\ [p, A;\ 1\text{-}p, C] \sim B$
Substitutability	$(A \sim B) \Rightarrow [p, A;\ 1\text{-}p, C] \sim [p, B;\ 1\text{-}p, C]$
Monotonicity	$(A \succ B) \Rightarrow (p \geq q \Leftrightarrow [p, A;\ 1\text{-}p, B] \succsim [q, A;\ 1\text{-}q, B])$
Decomposability	$[p, A;\ 1\text{-}p, [q, B;\ 1\text{-}q, C]] \sim [p, A;\ (1\text{-}p)q, B;\ (1\text{-}p)(1\text{-}q), C]$

It follows from these axioms that there exists a real-valued function U that operates on states such that

$$U(A) > U(B) \Leftrightarrow A \succ B$$
$$U(A) = U(B) \Leftrightarrow A \sim B$$

The St. Petersburg ”paradox”

You are offered to play the following game (bet): You flip a coin repeatedly until you get your first heads. You will then be paid $2 to the power of every flip you made, including the final one (the price matrix is below).

How much are you willing to pay to participate (participation is not free)?

Toss	**Winning**
H	$2
TH	$4
TTH	$8
TTTH	$16
...	...

The St. Petersburg "paradox"

What is the expected winning in this betting game?

$$\langle \text{Winning} \rangle_N = \sum_{k=1}^{N} P(k) W(k) = \sum_{k=1}^{N} \frac{2^k}{2^k} = \sum_{k=1}^{N} 1 = N$$

A rational player should be willing to pay any sum of money to participate...
...if $ = Utility

The students in previous years' classes have offered $4 or less on average...

The St. Petersburg ”paradox”

Bernoulli (1738): The utility of money is not money; it is more like log(Money).

$$\langle \text{Utility} \rangle_N = \sum_{k=1}^{N} P(k)U(k) = \sum_{k=1}^{N} \frac{\ln(2^k)}{2^k} = \ln(2)\sum_{k=1}^{N} \frac{k}{2^k}$$

$$= 2\ln(2)\left[1 - \left(\frac{1}{2}\right)^N - N\left(\frac{1}{2}\right)^{N+1}\right]$$

$$\lim_{N\to\infty} \langle \text{Utility} \rangle_N = 2\ln(2)$$

Bernoulli's utility curve

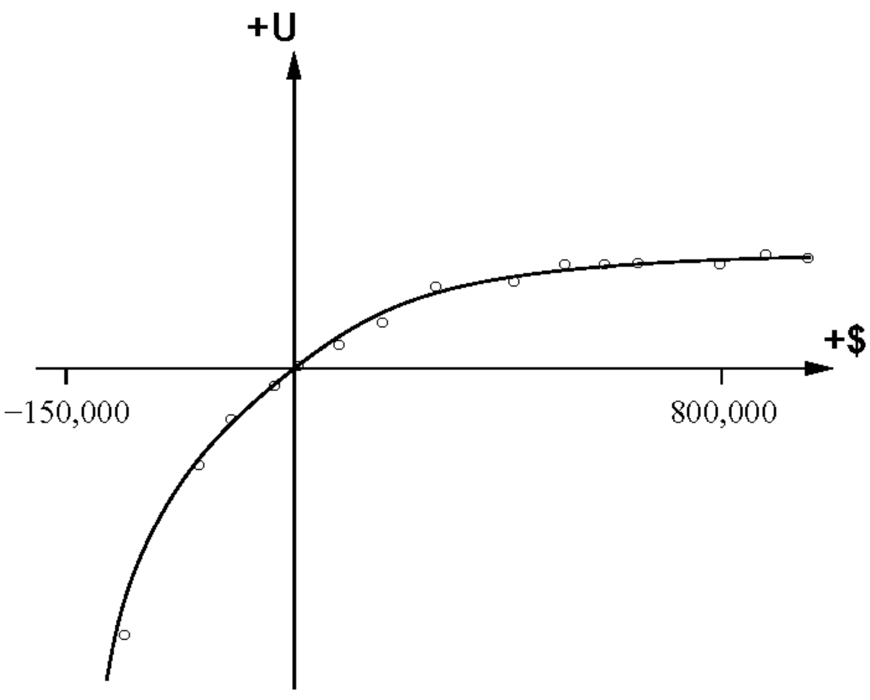

(a)

General ”human nature” utility curve

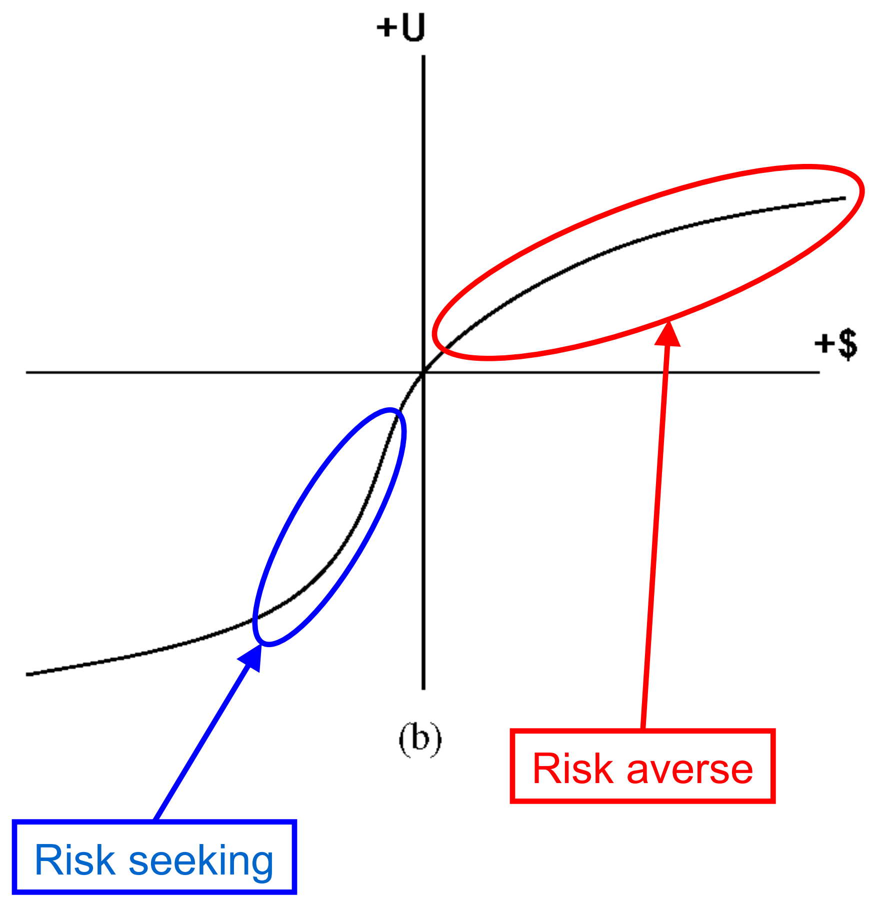

(b)

Lottery game 1

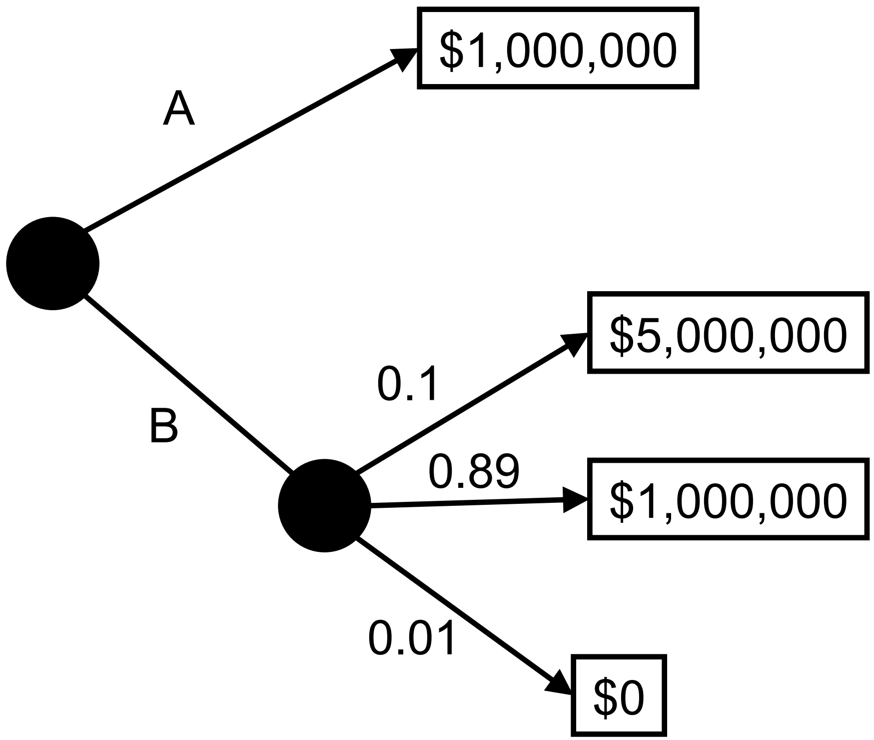

You can choose between alternatives A and B:

A) You get $1,000,000 for sure.

B) You can participate in a lottery where you can win up to $5 mill.

Lottery game 2

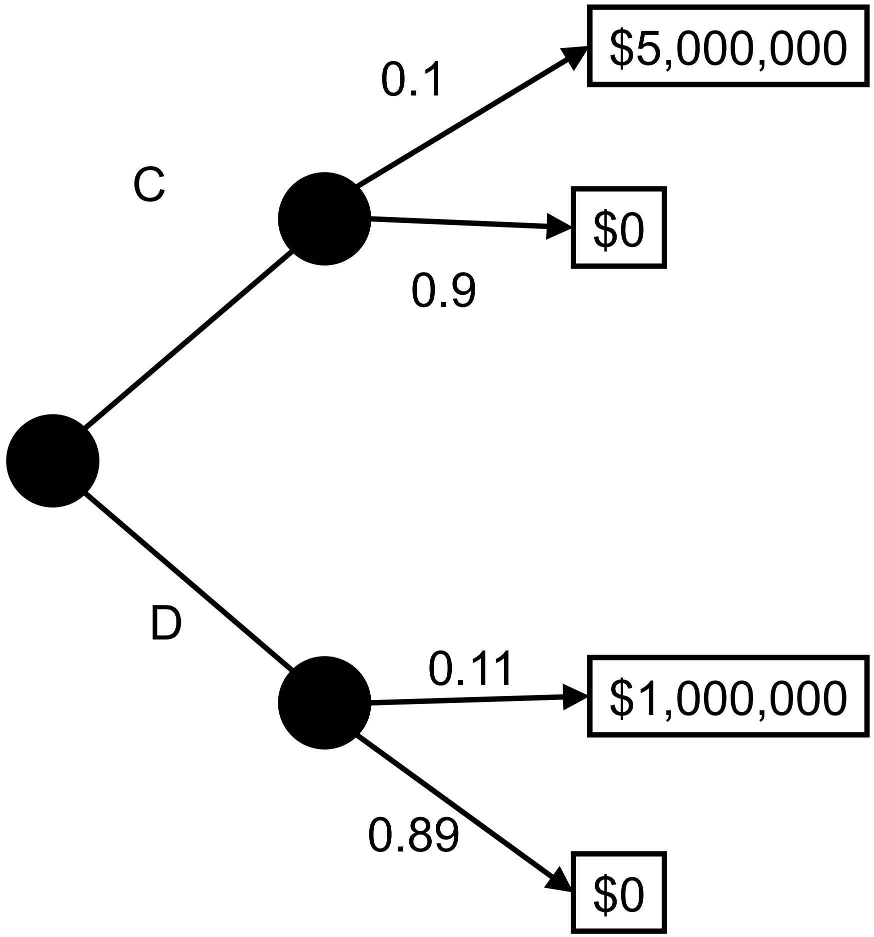

You can choose between alternatives C and D:

A) A lottery where you can win $5 mill.
B) A lottery where you can win $1 mill.

Lottery preferences

- People should select A and D, or B and C. Otherwise they are not being consistent...

$$U(A)-U(B)=U(\$1\text{M})-0.1U(\$5\text{M})-0.89U(\$1\text{M})-0.01U(\$0)=$$
$$0.11U(\$1\text{M})-0.1U(\$5\text{M})-0.01U(\$0)$$

$$U(D)-U(C)=0.11U(\$1\text{M})+0.89U(\$0)-0.1U(\$5\text{M})-0.9U(\$0)=$$
$$0.11U(\$1\text{M})-0.1U(\$5\text{M})-0.01U(\$0)$$

Allais paradox. Utility function does not capture a human's fear of looking like a complete idiot.

In last years classes, fewer than 50% have been consistent...

Form of *U(S)*

If the value of one attribute does not influence one's opinion about the preference for another attribute, then we have mutual preferential independence and can write:

$$V(X_1, X_2) = V_1(X_1) + V_2(X_2)$$

Where $V(X)$ is a value function (expressing the [monetary] value)

Example: The party problem

We are about to give a wedding party. It will be held during summer-time.

Should we be outdoors or indoors?

The party is such that we can't change our minds on the day of the party (different locations for indoors and outdoors).

What is the rational decision?

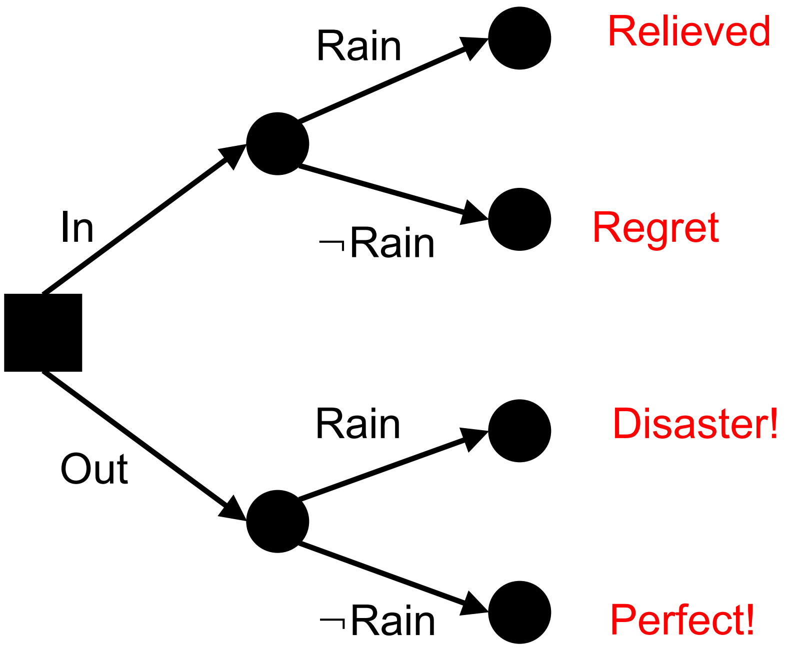

Example adapted from Breese & Kooler 1997

Example: The party problem

The value function: Assign a numerical (monetary) value to each outcome.

(We avoid the question on how this is done for the time being)

Location	Weather	Value
Indoors	Sun	$25
Indoors	Rain	$75
Outdoors	Sun	$100
Outdoors	Rain	$0

Let U(S) = log[V(S)+1]

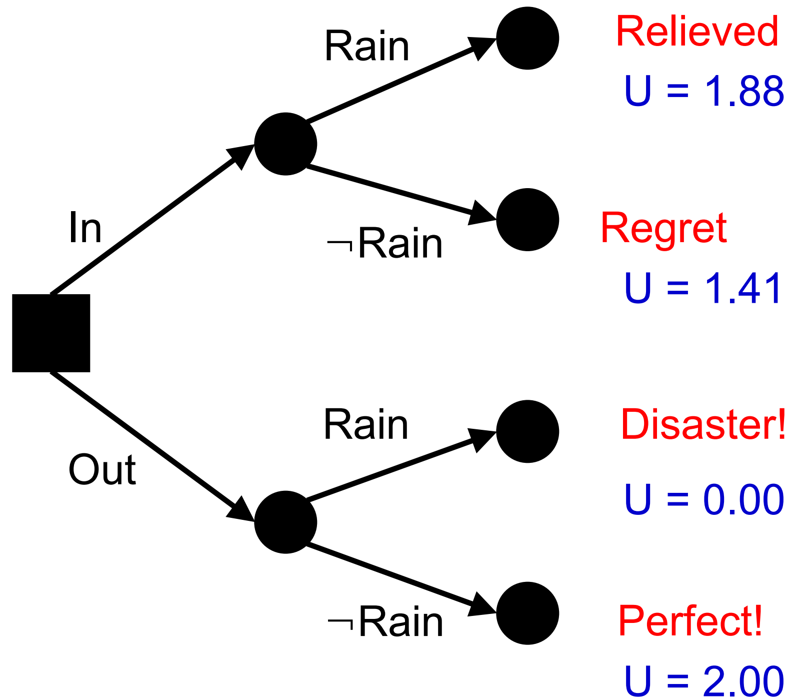

Example adapted from Breese & Kooler 1997

Example: The party problem

Get weather statistics for your location in the summer (June).

Location	P(Rain)
Stockholm, Sweden	18/30
Bergen, Norway	19/30
San Fransisco, USA	1/30
Seattle, USA	9/30
Paris, France	14/30
Haifa, Israel	0/30

Rain probabilities from Weatherbase
www.weatherbase.com/

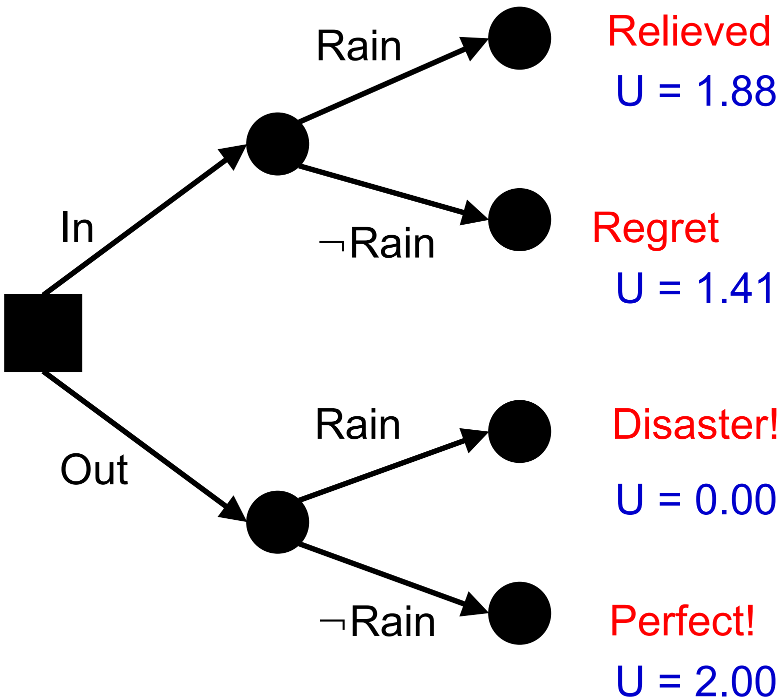

Example adapted from Breese & Kooler 1997

Example: The party problem

Example:
Stockholm, Sweden

$$EU(in) = \frac{18}{30} \times 1.88 + \frac{12}{30} \times 1.41 = 1.7$$

$$EU(out) = \frac{18}{30} \times 0.00 + \frac{12}{30} \times 2.00 = 0.8$$

Be indoors!

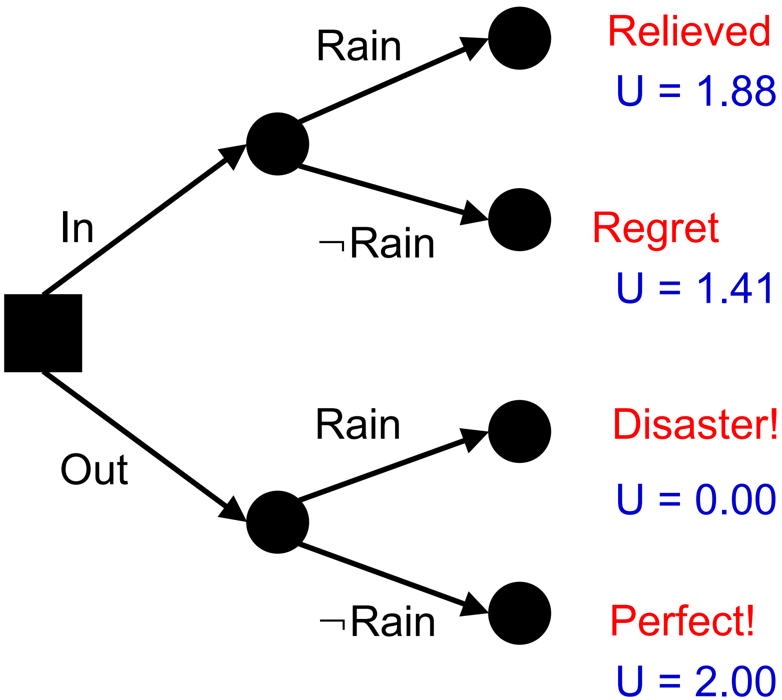

Example adapted from Breese & Kooler 1997

Example: The party problem

Example:
San Fransisco, California

$$EU(in) = \frac{1}{30} \times 1.88 + \frac{29}{30} \times 1.41 = 1.4$$

$$EU(out) = \frac{1}{30} \times 0.00 + \frac{29}{30} \times 2.00 = 1.9$$

Be outdoors!

The change from outdoors to indoors occurs at P(Rain) > 7/30

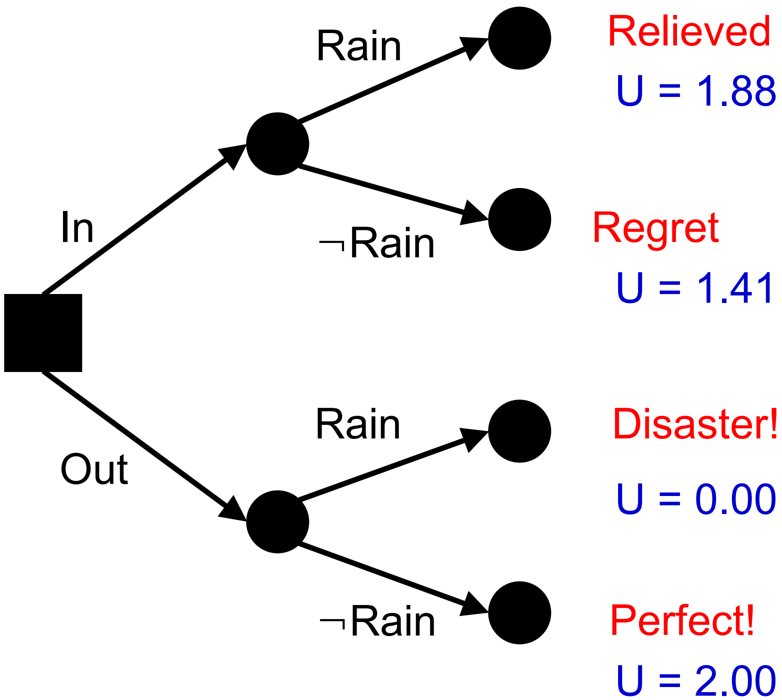

Example adapted from Breese & Kooler 1997

Decision network for the party problem

Decision represented by a rectangle

Chance (random variable) represented by an oval.

Utility function represented by a diamond.

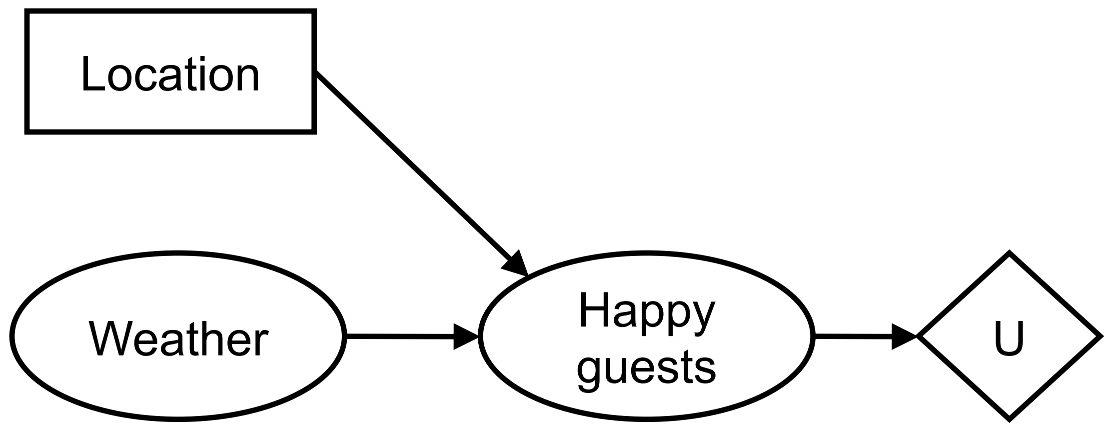

The value of information

- The value of a given piece of information is the difference in expected utility value between best actions before and after information is obtained.
- Information has value to the extent that it is likely to cause a change of plan and to the extent that the new plan will be significantly better than the old plan.