Lecture 4
Uncertainty

Marco Chiarandini

Department of Mathematics & Computer Science
University of Southern Denmark

Slides by Stuart Russell and Peter Norvig

Outline

Knowledge bases

Knowledge base = set of sentences in a **formal** language

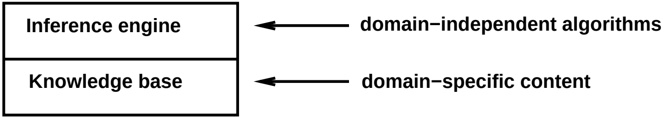

Declarative approach to building an agent (or other system):

Tell it what it needs to know

Then it can Ask itself what to do—answers should follow from the KB

Agents can be viewed at the knowledge level

i.e., **what they know**, regardless of how implemented

Or at the implementation level

i.e., data structures in KB and algorithms that manipulate them

A simple knowledge-based agent

```
function KB-Agent( percept) returns an action
   static: KB, a knowledge base
           t, a counter, initially 0, indicating time

   Tell(KB, Make-Percept-Sentence( percept, t))
   action ← Ask(KB, Make-Action-Query(t))
   Tell(KB, Make-Action-Sentence(action, t))
   t ← t + 1
   return action
```

The agent must be able to:

- Represent states, actions, etc.
- Incorporate new percepts
- Update internal representations of the world
- Deduce hidden properties of the world
- Deduce appropriate actions

Wumpus World PEAS description

Performance measure
gold +1000, death -1000
-1 per step, -10 for using the arrow

Environment
Squares adjacent to wumpus are smelly
Squares adjacent to pit are breezy
Glitter iff gold is in the same square
Shooting kills wumpus if you are facing it
Shooting uses up the only arrow
Grabbing picks up gold if in same square
Releasing drops the gold in same square

Actuators *LeftTurn*, *RightTurn*, *Forward*, *Grab*, *Release*, *Shoot*

Sensors *Breeze*, *Glitter*, *Smell*

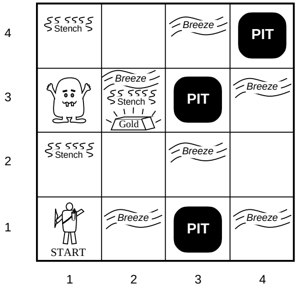

Wumpus world – Properties

Fully vs Partially observable??
No—only local perception
Deterministic vs Stochastic??
Deterministic—outcomes exactly specified
Episodic vs Sequential??
sequential at the level of actions
Static vs Dynamic??
Static—Wumpus and Pits do not move
Discrete vs Continous??
Discrete
Single-agent vs Multi-Agent??
Single—Wumpus is essentially a natural feature

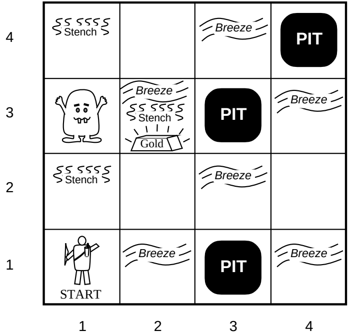

Exploring a wumpus world

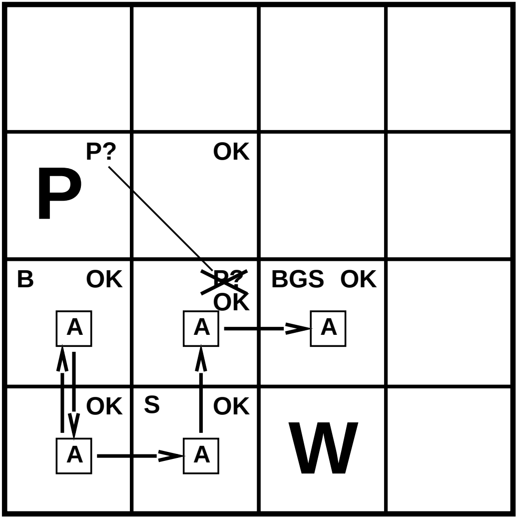

Other tight spots

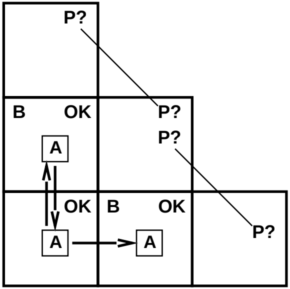

Breeze in (1,2) and (2,1)
$\Longrightarrow$ no safe actions

Assuming pits uniformly distributed,
(2,2) has pit w/ prob 0.86, vs. 0.31

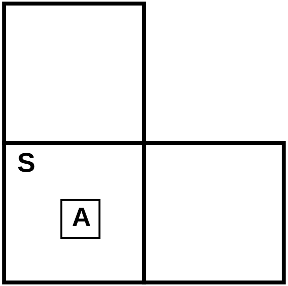

Smell in (1,1)
$\Longrightarrow$ cannot move
Can use a strategy of coercion:
shoot straight ahead
wumpus was there $\Longrightarrow$ dead $\Longrightarrow$ safe
wumpus wasn't there $\Longrightarrow$ safe

Outline

Logic in general

Logics are formal languages for representing information
such that conclusions can be drawn

Syntax defines the sentences in the language

Semantics define the "meaning" of sentences;
i.e., define truth of a sentence in a world

E.g., the language of arithmetic

$x + 2 \geq y$ is a sentence; $x2 + y >$ is not a sentence

$x + 2 \geq y$ is true iff the number $x + 2$ is no less than the number y

$x + 2 \geq y$ is true in a world where $x = 7,\ y = 1$

$x + 2 \geq y$ is false in a world where $x = 0,\ y = 6$

Entailment

Entailment means that one thing **follows from** another:

$$KB \models \alpha$$

Knowledge base KB entails sentence α
if and only if
α is true in all worlds where KB is true

E.g., the KB containing "OB won" and "FCK won"
entails "Either OB won or FCK won"

E.g., $x + y = 4$ entails $4 = x + y$

Entailment is a relationship between sentences (i.e., **syntax**)
that is based on **semantics**
Key idea: brains process **syntax** (of some sort)
trying to reproduce this mechanism

Models

Logicians typically think in terms of models, which are formally structured worlds with respect to which truth can be evaluated

We say m is a model of a sentence α if α is true in m

$M(\alpha)$ is the set of all models of α

Then $KB \models \alpha$ if and only if $M(KB) \subseteq M(\alpha)$

E.g. KB = OB won and FCK won
α = OB won

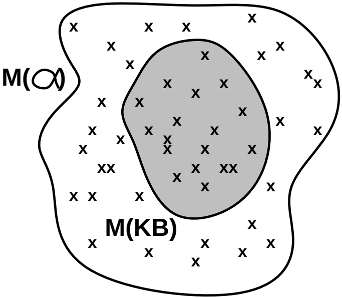

Entailment in the wumpus world

Situation after detecting nothing in [1,1], moving right, breeze in [2,1]

Consider possible models for ?s assuming only pits

3 Boolean choices $\Longrightarrow$ 8 possible models

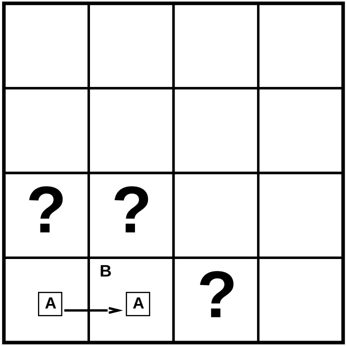

Wumpus models

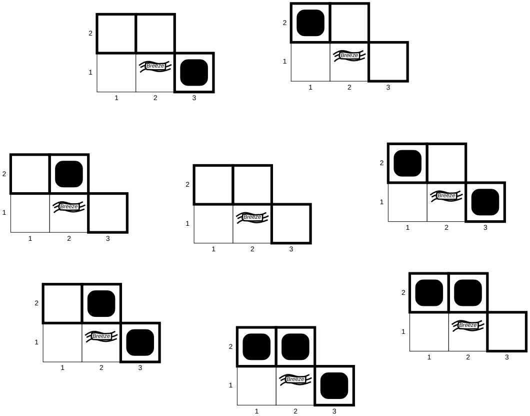

Wumpus models

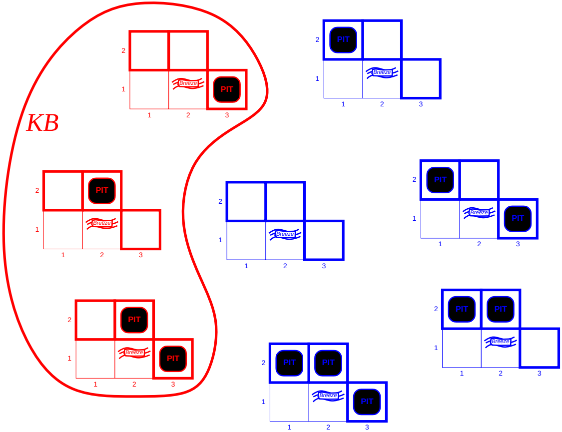

KB = wumpus-world rules + observations

Wumpus models

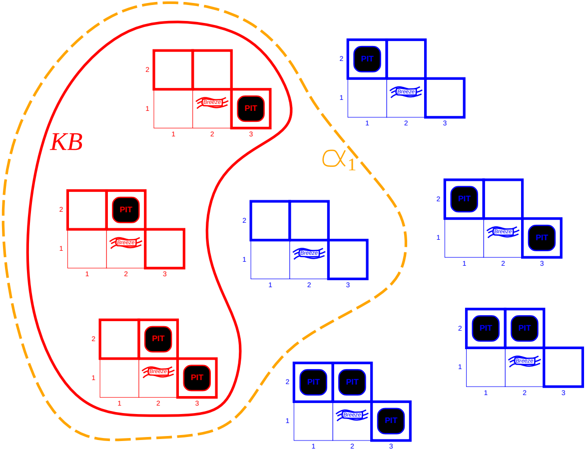

KB = wumpus-world rules + observations

α_1 = "[1,2] is safe", $KB \models \alpha_1$, proved by model checking

Wumpus models

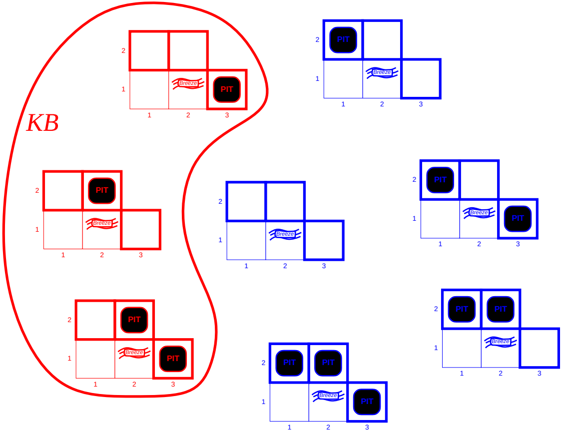

KB = wumpus-world rules + observations

Wumpus models

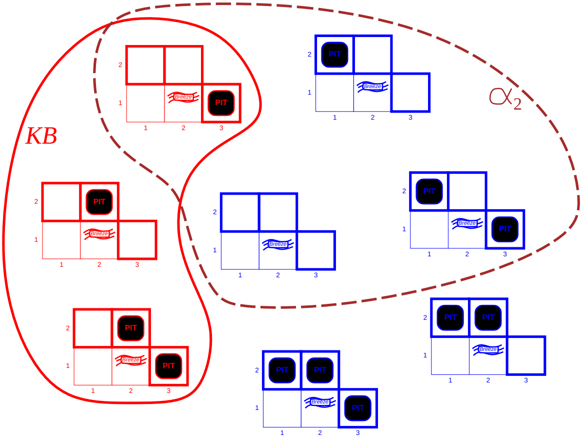

KB = wumpus-world rules + observations
α_2 = "[2,2] is safe", $KB \not\models \alpha_2$

Inference

$KB \vdash_i \alpha$ = sentence α can be derived from KB by procedure i

Consequences of KB are a haystack; α is a needle.
Entailment = needle in haystack; inference = finding it

Soundness: i is sound if
whenever $KB \vdash_i \alpha$, it is also true that $KB \models \alpha$

Completeness: i is complete if
whenever $KB \models \alpha$, it is also true that $KB \vdash_i \alpha$

Preview: we will define a logic (first-order logic) which is expressive enough to say almost anything of interest, and for which there exists a sound and complete inference procedure.

That is, the procedure will answer any question whose answer follows from what is known by the KB.

Outline

Outline

- ♢ Uncertainty
- ♢ Probability
- ♢ Syntax and Semantics
- ♢ Inference
- ♢ Independence and Bayes' Rule

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 (KCBS traffic reports)
3) uncertainty in action outcomes (flat tire, etc.)
4) immense complexity of modelling and predicting traffic

Hence a purely logical approach either

1. risks falsehood: "A_{25} will get me there on time"
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."

(A_{1440} might reasonably be said to get me there on time
but I'd have to stay overnight in the airport ...)

Methods for handling uncertainty

Logic-based abductive inference: Default or nonmonotonic logic:
Assume my car does not have a flat tire
Assume A_{25} works unless contradicted by evidence
Issues: What assumptions are reasonable? How to handle contradiction?

Rules with fudge factors:
$A_{25} \mapsto_{0.3} AtAirportOnTime$
$Sprinkler \mapsto_{0.99} WetGrass$
$WetGrass \mapsto_{0.7} Rain$
Issues: Problems with combination, e.g., *Sprinkler* causes *Rain*??

Probability
Given the available evidence,
A_{25} will get me there on time with probability 0.04
Mahaviracarya (9th C.), Cardano (1565) theory of gambling

(Fuzzy logic handles **degree of truth** NOT uncertainty e.g.,
WetGrass is true to degree 0.2)

Probability

Probabilistic assertions **summarize** effects of
laziness: failure to enumerate exceptions, qualifications, etc.
ignorance: lack of relevant facts, initial conditions, etc. inherent stochasticity: toss of coin, roll of a dice, 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., $P(A_{25}|\text{no reported accidents, 5 a.m.}) = 0.15$
(Analogous to logical entailment status $KB \models \alpha$, not truth.)

Making decisions under uncertainty

Suppose I believe the following:

$$
\begin{aligned}
P(A_{25} \text{ gets me there on time}|\ldots) &= 0.04 \\
P(A_{90} \text{ gets me there on time}|\ldots) &= 0.70 \\
P(A_{120} \text{ gets me there on time}|\ldots) &= 0.95 \\
P(A_{1440} \text{ gets me there on time}|\ldots) &= 0.9999
\end{aligned}
$$

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

Interpretations of Probability

- Classical interpretation: probabilities can be determined a priori by an examination of the space of possibilities.
  It assigns probabilities in the absence of any evidence, or in the presence of symmetrically balanced evidence

- Logical interpretation: generalizes the classcial it in two important ways:
  - possibilities may be assigned unequal weights
  - probabilities can be computed whatever the evidence may be, symmetrically balanced or not

- Frequentist: the probability of an attribute A in a finite reference class B is the relative frequency of actual occurrences of A within B.
  issue of identity

- Propensity interpretation: innate property of the objects

- Subjective interpretation: subjective degree of belief + betting system to avoid unconstrained subjectivism

Probability basics

DEFINITION

ELEMENTARY EVENT

An *elementary* or *atomic event* is a happening or occurrence that cannot be made up of other events.

EVENT, E

An *event* is a set of elementary events.

SAMPLE SPACE, S

The set of all possible outcomes of an event E is the *sample space* S or *universe* for that event.

PROBABILITY, p

The *probability* of an event E in a sample space S is the ratio of the number of elements in E to the total number of possible outcomes of the sample space S of E. Thus, $p(E) = |E| / |S|$.

Probability basics

The probability of any event E from the sample space S is:

$0 \leq p(E) \leq 1$, where $E \subseteq S$

The sum of the probabilities of all possible outcomes is 1

The probability of the complement of an event is

$p(\overline{E}) = (|S| - |E|) \; / \; |S| = (|S| \; / \; |S|) - (|E| \; / \; |S|) = 1 - p(E).$

The probability of the contradictory or false outcome of an event

$p(\{\ \}) = 1 - p(\overline{\{\ \}}) = 1 - p(S) = 1 - 1 = 0$, or alternatively,

$= |\{\ \}| \; / \; |S| = 0 \; / \; |S| = 0$

Probability basics

DEFINITION

INDEPENDENT EVENTS

Two events A and B are *independent* if and only if the probability of their both occurring is equal to the product of their occurring individually. This independence relation is expressed:

$$p(A \cap B) = p(A) * p(B)$$

We sometimes use the equivalent notation p(s,d) for $p(s \cap d)$. We clarify the notion of independence further in the context of conditional probabilities in Section 5.2.4.

The three Kolmogorov Axioms

1. The probability of event E in sample space S is between 0 and 1, ie, $0 \leq p(E) \leq 1$
2. When the union of all E gives S, $p(S) = 1$ and $p(\bar{S}) = 0$
3. The probability of the union of two sets of events A and B is:

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

Probability basics

DEFINITION

RANDOM VARIABLE

A *random variable* is a function whose domain is a sample space and range a set of outcomes, most often real numbers. Rather than using a problem-specific event space, a random variable allows us to talk about probabilities as numerical values that are related to an event space.

BOOLEAN, DISCRETE, and CONTINUOUS RANDOM VARIABLES

A *boolean random variable* is a function from an event space to {true, false} or to the subset of real numbers {0.0, 1.0}. A boolean random variable is sometimes called a *Bernoulli trial*.

A *discrete random variable*, which includes boolean random variables as a subset, is a function from the sample space to (a countable subset of) real numbers in [0.0, 1.0].

A *continuous random variable* has as its range the set of real numbers.

Propositions

Think of a proposition as the event (set of sample points) where the proposition is true

Given Boolean random variables A and B:
event a = set of sample points where $A = true$
event $\neg a$ = set of sample points where $A = false$
event $a \wedge b$ = points where $A = true$ and $B = true$

Often in AI applications, the sample points are **defined** by the values of a set of random variables, i.e., the sample space is the Cartesian product of the ranges of the variables

With Boolean variables, sample point = propositional logic model
e.g., $A = true$, $B = false$, or $a \wedge \neg b$.
Proposition = disjunction of atomic events in which it is true
e.g., $(a \vee b) \equiv (\neg a \wedge b) \vee (a \wedge \neg b) \vee (a \wedge b)$
$\implies P(a \vee b) = P(\neg a \wedge b) + P(a \wedge \neg b) + P(a \wedge b)$

Why use probability?

The definitions imply that certain logically related events must have related probabilities
E.g., $P(a \vee b) = P(a) + P(b) - P(a \wedge b)$

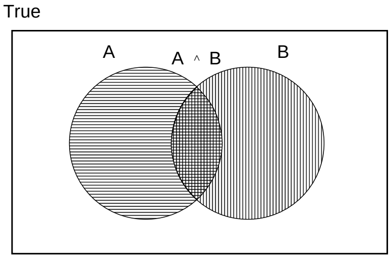

de Finetti (1931): an agent who bets according to probabilities that violate these axioms can be forced to bet so as to lose money regardless of outcome.

Syntax for propositions

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 $\langle sunny, rain, cloudy, snow \rangle$
$Weather = rain$ is a proposition
Values must be exhaustive and mutually exclusive

Continuous random variables (bounded or unbounded)
e.g., $Temp = 21.6$; also allow, e.g., $Temp < 22.0$.

Prior probability

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

Probability distribution gives values for all possible assignments:
$\Pr(Weather) = \langle 0.72, 0.1, 0.08, 0.1 \rangle$ (normalized, i.e., sums to 1)

Joint probability distribution for a set of r.v.s gives the probability of every atomic event on those r.v.s (i.e., every sample point)
$\Pr(Weather, Cavity) =$ a 4×2 matrix of values:

$Weather =$	$sunny$	$rain$	$cloudy$	$snow$
$Cavity = true$	0.144	0.02	0.016	0.02
$Cavity = false$	0.576	0.08	0.064	0.08

Every question about a domain can be answered by the joint distribution because every event is a sum of sample points

Probability for continuous variables

Express distribution as a parameterized function of value:

$P(X = x) = U[18, 26](x)$ = uniform density between 18 and 26

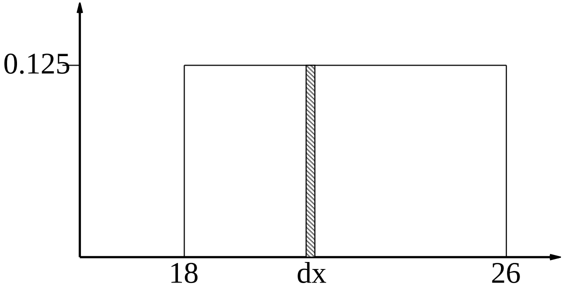

Here P is a density; integrates to 1.

$P(X = 20.5) = 0.125$ really means

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

Gaussian density

$$P(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2}$$

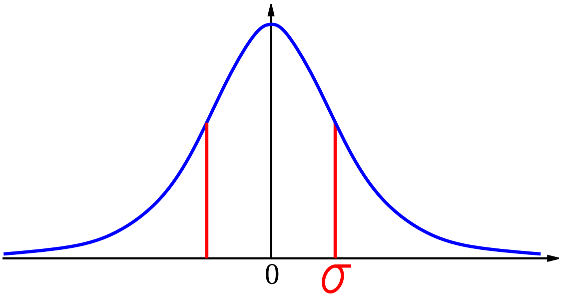

Rules to remember

- Complementarity $\Pr(B) = 1 - \Pr(A)$
- Marginalization $\Pr(B) = \sum_a \Pr(B, A = a)$
- Total probability $\Pr(B) = \sum_a \Pr(B|A = a) \Pr(A = a)$
- Dependency or Conditional probability $\Pr(A \mid B) = \frac{\Pr(A,B)}{\Pr(B)}$
- Product rule $\Pr(A, B) = \Pr(A) \Pr(B)$
- Normalization $\Pr(A \mid e) = \alpha \Pr(A, e)$
- Chain rule

$$\begin{aligned} \Pr(A_1 \cup A_2 \cup \cdots \cup A_n) &= \Pr(A_1) \Pr(A_2 \mid A_1, A_2) \ldots \Pr(A_n \mid A_{n-1}, A_{n-2} \ldots A_1) \\ &= \prod_{i=1}^{n} (A_i \mid A_{i-1}, A_{i-2}, \ldots, A_1) \end{aligned}$$

- Bayes' rule $\Pr(C \mid E) = \frac{\Pr(E|C) \Pr(C)}{\Pr(E)} = \alpha \Pr(E \mid C) \Pr(C)$
- Conditional Independence $\Pr(E_1, E_2 \mid C) = \Pr(E_1 \mid C) \Pr(E_2 \mid C)$ or $\Pr(E_1 \mid C, E_2) = \Pr(E_1 \mid C)$

Conditional probability

Conditional or posterior probabilities

e.g., $P(cavity|toothache) = 0.8$

i.e., **given that** *toothache* **is all I know**

NOT "if *toothache* then 80% chance of *cavity*" (Notation for conditional distributions: $\Pr(Cavity|Toothache)$ = 2-element vector of 2-element vectors)

If we know more, e.g., *cavity* is also given, then we have

$P(cavity|toothache, cavity) = 1$

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.,

$P(cavity|toothache, 49ersWin) = P(cavity|toothache) = 0.8$

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

Conditional probability

Definition

Conditional probability:

$$P(a|b) = \frac{P(a \wedge b)}{P(b)} \text{ if } P(b) \neq 0$$

Product rule gives an alternative formulation:
$P(a \wedge b) = P(a|b)P(b) = P(b|a)P(a)$

A general version holds for whole distributions, e.g.,
$\Pr(Weather, Cavity) = \Pr(Weather|Cavity)\Pr(Cavity)$
(View as a 4×2 set of equations, **not** matrix mult.)

Definition

Chain rule is derived by successive application of product rule:

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

Inference by enumeration

Start with the joint distribution:

	toothache		¬ *toothache*	
	catch	¬ *catch*	*catch*	¬ *catch*
cavity	**.108**	**.012**	**.072**	**.008**
¬ *cavity*	**.016**	**.064**	**.144**	**.576**

For any proposition ϕ, sum the atomic events where it is true:

$$P(\phi) = \sum_{\omega:\omega\models\phi} P(\omega)$$

Inference by enumeration

Start with the joint distribution:

	toothache		¬ *toothache*	
	catch	¬ *catch*	*catch*	¬ *catch*
cavity	**.108**	**.012**	**.072**	**.008**
¬ *cavity*	**.016**	**.064**	**.144**	**.576**

For any proposition ϕ, sum the atomic events where it is true:

$$P(\phi) = \sum_{\omega:\omega \models \phi} P(\omega)$$

$$P(toothache) = 0.108 + 0.012 + 0.016 + 0.064 = 0.2$$

Inference by enumeration

Start with the joint distribution:

	toothache		*¬ toothache*	
	catch	*¬ catch*	*catch*	*¬ catch*
cavity	**.108**	**.012**	**.072**	**.008**
¬ cavity	**.016**	**.064**	**.144**	**.576**

For any proposition ϕ, sum the atomic events where it is true:

$$P(\phi) = \sum_{\omega:\omega\models\phi} P(\omega)$$

$$P(cavity \vee toothache) = 0.108 + 0.012 + 0.072 + 0.008 + 0.016 + 0.064 = 0.28$$

Inference by enumeration

Start with the joint distribution:

	toothache		¬ *toothache*	
	catch	¬ *catch*	*catch*	¬ *catch*
cavity	**.108**	**.012**	**.072**	**.008**
¬ *cavity*	**.016**	**.064**	**.144**	**.576**

Can also compute conditional probabilities:

$$P(\neg cavity | toothache) = \frac{P(\neg cavity \land toothache)}{P(toothache)}$$

$$= \frac{0.016 + 0.064}{0.108 + 0.012 + 0.016 + 0.064} = 0.4$$

Normalization

	toothache		¬ *toothache*	
	catch	¬ *catch*	*catch*	¬ *catch*
cavity	**.108**	**.012**	**.072**	**.008**
¬ *cavity*	**.016**	**.064**	**.144**	**.576**

Denominator can be viewed as a normalization constant α

$$\mathrm{Pr}(Cavity|toothache) = \alpha\, \mathrm{Pr}(Cavity, toothache)$$

$$= \alpha\,[\mathrm{Pr}(Cavity, toothache, catch) + \mathrm{Pr}(Cavity, toothache, \neg catch)]$$

$$= \alpha\,[\langle 0.108, 0.016\rangle + \langle 0.012, 0.064\rangle]$$

$$= \alpha\,\langle 0.12, 0.08\rangle = \langle 0.6, 0.4\rangle$$

General idea: compute distribution on query variable by fixing evidence variables and summing over hidden variables

Inference by enumeration, contd.

Let $\mathbf{X}$ be all the variables. Typically, we want
the posterior joint distribution of the query variables $\mathbf{Y}$
given specific values $\mathbf{e}$ for the evidence variables $\mathbf{E}$
Let the hidden variables be $\mathbf{H} = \mathbf{X} - \mathbf{Y} - \mathbf{E}$

Then the required summation of joint entries is done by summing out the hidden variables:

$$\Pr(\mathbf{Y}|\mathbf{E}=\mathbf{e}) = \alpha \Pr(\mathbf{Y}, \mathbf{E}=\mathbf{e}) = \alpha \sum_{\mathbf{h}} \Pr(\mathbf{Y}, \mathbf{E}=\mathbf{e}, \mathbf{H}=\mathbf{h})$$

The terms in the summation are joint entries because $\mathbf{Y}$, $\mathbf{E}$, and $\mathbf{H}$ together exhaust the set of random variables

Obvious problems:
1) Worst-case time complexity $O(d^n)$ where d is the largest arity
2) Space complexity $O(d^n)$ to store the joint distribution
3) How to find the numbers for $O(d^n)$ entries???

Summary

- Interpretations of probability
- Axioms of Probability
- (Continuous/Discrete) Random Variables
- Prior probability, joint probability, conditional or posterior probability, chain rule
- Inference by enumeration

How to reduce the computation of inference?

Probability basics

DEFINITION

INDEPENDENT EVENTS

Two events A and B are *independent* of each other if and only if $p(A \cap B) = p(A)\, p(B)$. When $p(B) \neq 0$ this is the same as saying that $p(A) = p(A|B)$. That is, knowing that B is true does not affect the probability of A being true.

CONDITIONALLY INDEPENDENT EVENTS

Two events A and B are said to be *conditionally independent* of each other, given event C if and only if $p((A \cap B) \mid C) = p(A \mid C)\, p(B \mid C)$.

Independence

A and B are independent iff
$\Pr(A \mid B) = \Pr(A)$ or $\Pr(B \mid A) = \Pr(B)$ or $\Pr(A, B) = \Pr(A)\Pr(B)$

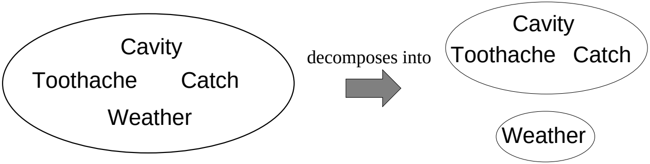

$$\Pr(Toothache, Catch, Cavity, Weather)$$
$$= \Pr(Toothache, Catch, Cavity)\Pr(Weather)$$

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

Absolute independence powerful but rare

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

Conditional independence

$\Pr(\mathit{Toothache}, \mathit{Cavity}, \mathit{Catch})$ 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) $P(\mathit{catch} \mid \mathit{toothache}, \mathit{cavity}) = P(\mathit{catch} \mid \mathit{cavity})$

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

(2) $P(\mathit{catch} \mid \mathit{toothache}, \neg\mathit{cavity}) = P(\mathit{catch} \mid \neg\mathit{cavity})$

Catch is conditionally independent of *Toothache* given *Cavity*:

$\Pr(\mathit{Catch} \mid \mathit{Toothache}, \mathit{Cavity}) = \Pr(\mathit{Catch} \mid \mathit{Cavity})$

Equivalent statements:

$\Pr(\mathit{Toothache} \mid \mathit{Catch}, \mathit{Cavity}) = \Pr(\mathit{Toothache} \mid \mathit{Cavity})$

$\Pr(\mathit{Toothache}, \mathit{Catch} \mid \mathit{Cavity}) = \Pr(\mathit{Toothache} \mid \mathit{Cavity}) \Pr(\mathit{Catch} \mid \mathit{Cavity})$

Conditional independence contd.

Write out full joint distribution using chain rule:

$$\begin{aligned}
&\Pr(\mathit{Toothache}, \mathit{Catch}, \mathit{Cavity}) \\
&= \Pr(\mathit{Toothache} \mid \mathit{Catch}, \mathit{Cavity}) \Pr(\mathit{Catch}, \mathit{Cavity}) \\
&= \Pr(\mathit{Toothache} \mid \mathit{Catch}, \mathit{Cavity}) \Pr(\mathit{Catch} \mid \mathit{Cavity}) \Pr(\mathit{Cavity}) \\
&= \Pr(\mathit{Toothache} \mid \mathit{Cavity}) \Pr(\mathit{Catch} \mid \mathit{Cavity}) \Pr(\mathit{Cavity})
\end{aligned}$$

I.e., $2 + 2 + 1 = 5$ 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

Product rule $P(a \wedge b) = P(a \mid b)P(b) = P(b \mid a)P(a)$

$$\implies \text{Bayes' rule } P(a \mid b) = \frac{P(b \mid a)P(a)}{P(b)}$$

or in distribution form

$$\Pr(Y \mid X) = \frac{\Pr(X \mid Y)\Pr(Y)}{\Pr(X)} = \alpha \Pr(X \mid Y)\Pr(Y)$$

Useful for assessing diagnostic probability from causal probability:

$$P(Cause \mid Effect) = \frac{P(Effect \mid Cause)P(Cause)}{P(Effect)}$$

E.g., let M be meningitis, S be stiff neck:

$$P(m \mid s) = \frac{P(s \mid m)P(m)}{P(s)} = \frac{0.8 \times 0.0001}{0.1} = 0.0008$$

Note: posterior probability of meningitis still very small!

Outline

Wumpus World

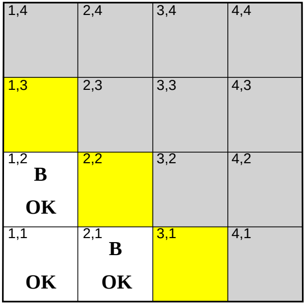

$P_{ij} = true$ iff $[i, j]$ contains a pit
$B_{ij} = true$ iff $[i, j]$ is breezy
Include only $B_{1,1}, B_{1,2}, B_{2,1}$ in the probability model

Specifying the probability model

The full joint distribution is $\Pr(P_{1,1}, \ldots, P_{4,4}, B_{1,1}, B_{1,2}, B_{2,1})$

Apply product rule: $\Pr(B_{1,1}, B_{1,2}, B_{2,1} \mid P_{1,1}, \ldots, P_{4,4}) \Pr(P_{1,1}, \ldots, P_{4,4})$

(Do it this way to get $P(\mathit{Effect} \mid \mathit{Cause})$.)

First term: 1 if pits are adjacent to breezes, 0 otherwise

Second term: pits are placed randomly, probability 0.2 per square:

$$\Pr(P_{1,1}, \ldots, P_{4,4}) = \prod_{i,j=1,1}^{4,4} \Pr(P_{i,j}) = 0.2^n \times 0.8^{16-n}$$

for n pits.

Observations and query

We know the following facts:

$b = \neg b_{1,1} \wedge b_{1,2} \wedge b_{2,1}$

$known = \neg p_{1,1} \wedge \neg p_{1,2} \wedge \neg p_{2,1}$

Query is $\Pr(P_{1,3} \mid known, b)$

Define $Unknown$ = P_{ij}s other than $P_{1,3}$ and $Known$

For inference by enumeration, we have

$$\Pr(P_{1,3} \mid known, b) = \alpha \sum_{unknown} \Pr(P_{1,3}, unknown, known, b)$$

Grows exponentially with number of squares!

Using conditional independence

Basic insight: observations are conditionally independent of other hidden squares given neighbouring hidden squares

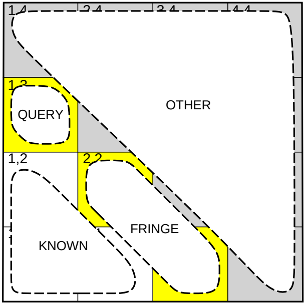

Define $Unknown = Fringe \cup Other$

$\Pr(b \mid P_{1,3}, Known, Unknown) = \Pr(b \mid P_{1,3}, Known, Fringe)$

Manipulate query into a form where we can use this!

Using conditional independence contd.

$$\begin{aligned}
\Pr(P_{1,3} \mid known, b) &= \alpha \sum_{unknown} \Pr(P_{1,3}, unknown, known, b) \\
&= \alpha \sum_{unknown} \Pr(b \mid P_{1,3}, known, unknown) \Pr(P_{1,3}, known, unknown) \\
&= \alpha \sum_{fringe} \sum_{other} \Pr(b \mid known, P_{1,3}, fringe, other) \Pr(P_{1,3}, known, fringe, other) \\
&= \alpha \sum_{fringe} \sum_{other} \Pr(b \mid known, P_{1,3}, fringe) \Pr(P_{1,3}, known, fringe, other) \\
&= \alpha \sum_{fringe} \Pr(b \mid known, P_{1,3}, fringe) \sum_{other} \Pr(P_{1,3}, known, fringe, other) \\
&= \alpha \sum_{fringe} \Pr(b \mid known, P_{1,3}, fringe) \sum_{other} \Pr(P_{1,3}) P(known) P(fringe) P(other) \\
&= \alpha\, P(known) \Pr(P_{1,3}) \sum_{fringe} \Pr(b \mid known, P_{1,3}, fringe) P(fringe) \sum_{other} P(other) \\
&= \alpha' \Pr(P_{1,3}) \sum_{fringe} \Pr(b \mid known, P_{1,3}, fringe) P(fringe)
\end{aligned}$$

Using conditional independence contd.

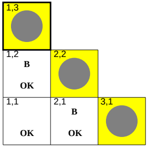

0.2 x 0.2 = 0.04

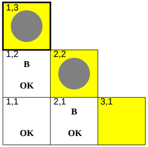

0.2 x 0.8 = 0.16

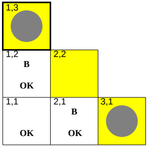

0.8 x 0.2 = 0.16

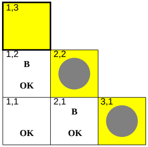

0.2 x 0.2 = 0.04

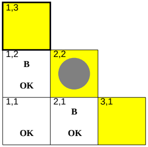

0.2 x 0.8 = 0.16

$$\begin{aligned}\Pr(P_{1,3} \mid known, b) &= \alpha' \langle 0.2(0.04 + 0.16 + 0.16),\ 0.8(0.04 + 0.16) \rangle \\ &\approx \langle 0.31, 0.69 \rangle\end{aligned}$$

$$\Pr(P_{2,2} \mid known, b) \approx \langle 0.86, 0.14 \rangle$$

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