

Introduction to Artificial Intelligence

Uncertainty

Chapter 14

Dieter Fox

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 (KUOV 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”
- 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.”

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

Outline

- ◇ Uncertainty
- ◇ Probability
- ◇ Syntax
- ◇ Semantics
- ◇ Inference rules

Methods for handling uncertainty

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:

$A_{25} \mapsto_{0.3}$ get there on time
 $Sprinkler \mapsto_{0.99} WetGrass$
 $WetGrass \mapsto_{0.7} Rain$

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.

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** assertions about the world

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:

$$P(A_{25} \text{ gets me there on time} | \dots) = 0.04$$

$$P(A_{90} \text{ gets me there on time} | \dots) = 0.70$$

$$P(A_{120} \text{ gets me there on time} | \dots) = 0.95$$

$$P(A_{1440} \text{ gets me there on time} | \dots) = 0.9999$$

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

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?

Axioms of probability


For any propositions A, B

$$1. 0 \leq P(A) \leq 1$$

$$2. P(True) = 1 \text{ and } P(False) = 0$$

$$3. P(A \vee B) = P(A) + P(B) - P(A \wedge B)$$

True



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

Similar to propositional logic: possible worlds defined by assignment of values to **random variables**.

Propositional or **Boolean** random variables

e.g., $Cavity$ (do I have a cavity?)

Include propositional logic expressions

e.g., $\neg Burglary \vee Earthquake$

Multivalued random variables

e.g., $Weather$ is one of $\langle sunny, rain, cloudy, snow \rangle$

Values must be exhaustive and mutually exclusive

Proposition constructed by assignment of a value:

e.g., $Weather = sunny$, also $Cavity = true$ for clarity

Syntax contd.

Conditional or **posterior probabilities**

e.g., $P(Cavity|Toothache) = 0.8$

i.e., **given that $Toothache$ is all I know**

Notation for conditional distributions:

$P(Weather|Earthquake) = 2\text{-element vector of 4-element vectors}$

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

Syntax contd.

Prior or **unconditional probabilities** of propositions

e.g., $P(Cavity) = 0.1$ and $P(Weather = sunny) = 0.72$

correspond to belief prior to arrival of any (new) evidence

Probability distribution gives values for all possible assignments:

$P(Weather) = \langle 0.72, 0.1, 0.08, 0.1 \rangle$ (**normalized**, i.e., sums to 1)

Joint probability distribution for a set of variables gives

values for each possible assignment to all the variables

$P(Weather, Cavity) =$ a 4×2 matrix of values:

$Weather =$	$sunny$	$rain$	$cloudy$	$snow$
$Cavity = true$				
$Cavity = false$				

Conditional probability

Definition of 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.,

$$P(Weather, Cavity) = P(Weather|Cavity)P(Cavity)$$

(View as a 4×2 set of equations, *not* matrix mult.)

Chain rule is derived by successive application of product rule:

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

Bayes' Rule

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

$$\Rightarrow \text{Bayes' rule } P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Why is this useful???

For assessing **diagnostic** probability from **causal** probability:

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

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

$$P(M|S) = \frac{P(S|M)P(M)}{P(S)} = \frac{0.8 \times 0.0001}{0.1} = 0.0008$$

Note: posterior probability of meningitis still very small!

Conditioning

Introducing a variable as an extra condition:

$$P(X|Y) = \sum_z P(X|Y, Z = z)P(Z = z|Y)$$

Intuition: often easier to assess each specific circumstance, e.g.,

$$\begin{aligned} P(RunOver|Cross) &= P(RunOver|Cross, Light = green)P(Light = green|Cross) \\ &\quad + P(RunOver|Cross, Light = yellow)P(Light = yellow|Cross) \\ &\quad + P(RunOver|Cross, Light = red)P(Light = red|Cross) \end{aligned}$$

When Y is absent, we have **summing out** or **marginalization**:

$$P(X) = \sum_z P(X|Z = z)P(Z = z) = \sum_z P(X, Z = z)$$

In general, given a joint distribution over a set of variables, the distribution over any subset (called a **marginal** distribution) can be calculated by summing out the other variables.

Normalization

Suppose we wish to compute a posterior distribution over A given $B = b$, and suppose A has possible values $a_1 \dots a_m$

We can apply Bayes' rule for each value of A :

$$P(A = a_1|B = b) = P(B = b|A = a_1)P(A = a_1)/P(B = b)$$

$$\vdots$$

$$P(A = a_m|B = b) = P(B = b|A = a_m)P(A = a_m)/P(B = b)$$

Adding these up, and noting that $\sum_i P(A = a_i|B = b) = 1$:

$$1/P(B = b) = 1/\sum_i P(B = b|A = a_i)P(A = a_i)$$

This is the **normalization factor**, constant w.r.t. i , denoted α :

$$P(A|B = b) = \alpha P(B = b|A)P(A)$$

Typically compute an unnormalized distribution, normalize at end

e.g., suppose $P(B = b|A)P(A) = \langle 0.4, 0.2, 0.2 \rangle$

then $P(A|B = b) = \alpha \langle 0.4, 0.2, 0.2 \rangle = \frac{\langle 0.4, 0.2, 0.2 \rangle}{0.4+0.2+0.2} = \langle 0.5, 0.25, 0.25 \rangle$

Full joint distributions

A **complete probability model** specifies every entry in the joint distribution for all the variables $\mathbf{X} = X_1, \dots, X_n$
I.e., a probability for each possible world $X_1 = x_1, \dots, X_n = x_n$

E.g., suppose *Toothache* and *Cavity* are the random variables:

	<i>Toothache</i> = <i>true</i>	<i>Toothache</i> = <i>false</i>
<i>Cavity</i> = <i>true</i>	0.04	0.06
<i>Cavity</i> = <i>false</i>	0.01	0.89

Possible worlds are mutually exclusive $\Rightarrow P(w_1 \wedge w_2) = 0$

Possible worlds are exhaustive $\Rightarrow w_1 \vee \dots \vee w_n$ is *True*

hence $\sum_i P(w_i) = 1$

Full joint distributions contd.

1) For any proposition ϕ defined on the random variables $\phi(w_i)$ is true or false

2) ϕ is equivalent to the disjunction of w_i s where $\phi(w_i)$ is true

$$\text{Hence } P(\phi) = \sum_{\{w_i: \phi(w_i)\}} P(w_i)$$

I.e., the unconditional probability of any proposition is computable as the sum of entries from the full joint distribution

Conditional probabilities can be computed in the same way as a ratio:

$$P(\phi|\xi) = \frac{P(\phi \wedge \xi)}{P(\xi)}$$

E.g.,

$$P(\text{Cavity}|\text{Toothache}) = \frac{P(\text{Cavity} \wedge \text{Toothache})}{P(\text{Toothache})} = \frac{0.04}{0.04 + 0.01} = 0.8$$