

# Uncertainty

## Chapter 13

### 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 (traffic reports)
3. uncertainty in action outcomes (flat tire, etc.)
4. immense complexity of modeling 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 ...)

## But...

- A decision must be made!
- No intelligent system can afford to consider all eventualities, wait until all the data is in and complete, or try all possibilities to see what happens

## Quick Overview of Reasoning Systems

- Logic: True or false, nothing in between. No uncertainty
- Probability: Degree of belief, but in the end it's either true or false

## Probability

Probabilistic assertions **summarize** effects of

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

**Subjective** probability:

- Probabilities relate propositions to agent's own state of knowledge  
e.g.,  $P(A_{25} \mid \text{no reported accidents}) = 0.06$

These are **not** assertions about the world

Probabilities of propositions change with new evidence:

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

# Making decisions under uncertainty

Suppose I believe the following:

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

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

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

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

- Which action to choose?

Depends on my **preferences** for missing flight vs. time spent waiting, etc.

- **Utility theory** is used to represent and infer preferences
- **Decision theory** = probability theory + utility theory

# Syntax

- Basic element: **random variable**
- Similar to propositional logic: possible worlds defined by assignment of values to random variables.
- **Boolean** random variables  
e.g., *Cavity* (do I have a cavity?)
- **Discrete** random variables  
e.g., *Weather* is one of *<sunny,rainy,cloudy,snow>*
- Domain values must be exhaustive and mutually exclusive
- Elementary proposition constructed by assignment of a value to a random variable: e.g., *Weather = sunny*, *Cavity = false* (abbreviated as  $\neg \text{cavity}$ )
- Complex propositions formed from elementary propositions and standard logical connectives e.g., *Weather = sunny*  $\vee$  *Cavity = false*

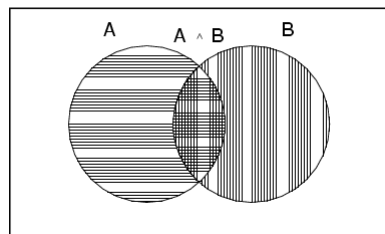
# Syntax

- **Atomic event:** A **complete** specification of the state of the world about which the agent is uncertain  
E.g., if the world consists of only two Boolean variables *Cavity* and *Toothache*, then there are 4 distinct atomic events:  
 $Cavity = false \wedge Toothache = false$   
 $Cavity = false \wedge Toothache = true$   
 $Cavity = true \wedge Toothache = false$   
 $Cavity = true \wedge Toothache = true$
- Atomic events are mutually exclusive and exhaustive

## Axioms of probability

- For any propositions  $A, B$ 
  - $0 \leq P(A) \leq 1$
  - $P(true) = 1$  and  $P(false) = 0$
  - $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$

True



# Prior probability

- **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
- **Probability distribution** gives values for all possible assignments:  
 $P(\text{Weather}) = \langle 0.72, 0.1, 0.08, 0.1 \rangle$  (**normalized**, i.e., sums to 1)
- **Joint probability distribution** for a set of random variables gives the probability of every atomic event on those random variables  
 $P(\text{Weather}, \text{Cavity})$  = a  $4 \times 2$  matrix of values:

<i>Weather</i> =	sunny	rainy	cloudy	snow
<i>Cavity</i> = true	0.144	0.02	0.016	0.02
<i>Cavity</i> = false	0.576	0.08	0.064	0.08

- Every question about a domain can be answered by the joint distribution

## How could we estimate the full joint distribution?

Parameter estimates are provided by expert knowledge, statistics on data samples, or a combination of both.

Suppose you have 20 variables.

Expert knowledge:

$P(X_1=0, X_2=0, \dots, X_{13}=1, \dots, X_{20}=0)$  vs.

$P(X_1=0, X_2=0, \dots, X_{13}=0, \dots, X_{20}=0)$  ?

Data Samples: practically speaking, we don't typically have enough data

## Conditional probability

- **Conditional or posterior probabilities**  
e.g.,  $P(\text{cavity} \mid \text{toothache}) = 0.8$   
i.e., given that *toothache* is all I know
- If we know more, e.g., *cavity* is also given, then we have  
 $P(\text{cavity} \mid \text{toothache}, \text{cavity}) = 1$
- New evidence may be irrelevant, allowing simplification,  
e.g.,  
 $P(\text{cavity} \mid \text{toothache}, \text{sunny}) = P(\text{cavity} \mid \text{toothache}) = 0.8$
- This kind of inference, sanctioned by domain knowledge,  
is crucial

## More on Conditional Probabilities

- $P(\text{CarWontStart} \mid \text{NoGas})$ 
  - This predicts a symptom based on an underlying cause
  - These can be generated empirically (Drain N gastanks, see how many cars start) or using expert knowledge
- $P(\text{NoGas} \mid \text{CarWontStart})$ 
  - Diagnosis. We have a symptom and want to predict the cause. This is what the system wants to determine

## Conditional probability

- Definition of conditional probability:  
 $P(a | b) = P(a \wedge b) / P(b)$  if  $P(b) > 0$
- **Product rule** gives an alternative formulation:  
 $P(a \wedge b) = P(a | b) P(b) = P(b | a) P(a)$
- **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  $P(a \wedge b) = P(a | b) P(b) = P(b | a) P(a)$   
 $\Rightarrow$  **Bayes' rule**:  $P(a | b) = P(b | a) P(a) / P(b)$
- Useful for assessing **diagnostic** probability from **causal** probability:  $\square$ 
  - $P(\text{Cause} | \text{Effect}) = P(\text{Effect} | \text{Cause}) P(\text{Cause}) / P(\text{Effect})$
  - E.g., let  $M$  be meningitis,  $S$  be stiff neck:  
 $P(m|s) = P(s|m) P(m) / P(s) = 0.8 \times 0.0001 / 0.1 = 0.0008$
  - Note: posterior probability of meningitis still very small!

## Inference by enumeration

- Start with the joint probability distribution:

	<i>toothache</i>		$\neg$ <i>toothache</i>	
	<i>catch</i>	$\neg$ <i>catch</i>	<i>catch</i>	$\neg$ <i>catch</i>
<i>cavity</i>	<b>.108</b>	<b>.012</b>	<b>.072</b>	<b>.008</b>
$\neg$ <i>cavity</i>	<b>.016</b>	<b>.064</b>	<b>.144</b>	<b>.576</b>

- For any proposition  $\phi$ , sum the atomic events where it is true:  $P(\phi) = \sum_{\omega: \omega \models \phi} P(\omega)$

## Inference by enumeration

- Start with the joint probability distribution:

	<i>toothache</i>		$\neg$ <i>toothache</i>	
	<i>catch</i>	$\neg$ <i>catch</i>	<i>catch</i>	$\neg$ <i>catch</i>
<i>cavity</i>	<b>.108</b>	<b>.012</b>	<b>.072</b>	<b>.008</b>
$\neg$ <i>cavity</i>	<b>.016</b>	<b>.064</b>	<b>.144</b>	<b>.576</b>

- For any proposition  $\phi$ , sum the atomic events where it is true:  $P(\phi) = \sum_{\omega: \omega \models \phi} P(\omega)$
- $P(\text{toothache}) = 0.108 + 0.012 + 0.016 + 0.064 = 0.2$

# Inference by enumeration

- Start with the joint probability distribution:

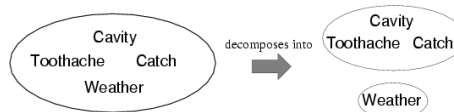
	<i>toothache</i>		$\neg$ <i>toothache</i>	
	<i>catch</i>	$\neg$ <i>catch</i>	<i>catch</i>	$\neg$ <i>catch</i>
<i>cavity</i>	.108	.012	.072	.008
$\neg$ <i>cavity</i>	.016	.064	.144	.576

- Can also compute conditional probabilities:

$$\begin{aligned}
 P(\neg \text{cavity} \mid \text{toothache}) &= \frac{P(\neg \text{cavity} \wedge \text{toothache})}{P(\text{toothache})} \\
 &= \frac{0.016 + 0.064}{0.108 + 0.012 + 0.016 + 0.064} \\
 &= 0.4
 \end{aligned}$$

# Independence

- $A$  and  $B$  are independent iff  
 $P(A/B) = P(A)$  or  $P(B/A) = P(B)$  or  $P(A, B) = P(A) P(B)$



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

- 32 entries reduced to 12; for  $n$  independent biased coins,  $O(2^n) \rightarrow O(n)$
- Absolute independence powerful but rare
- Dentistry is a large field with hundreds of variables, none of which are independent. What to do?

## Conditional independence

- If I have a cavity, the probability that the probe catches in it doesn't depend on whether I have a toothache:  
(1)  $P(\text{catch} \mid \text{toothache}, \text{cavity}) = P(\text{catch} \mid \text{cavity})$
- The same independence holds if I haven't got a cavity:  
(2)  $P(\text{catch} \mid \text{toothache}, \neg \text{cavity}) = P(\text{catch} \mid \neg \text{cavity})$   
□
- *Catch* is **conditionally independent** of *Toothache* given *Cavity*:  
 $P(\text{Catch} \mid \text{Toothache}, \text{Cavity}) = P(\text{Catch} \mid \text{Cavity})$

## Conditional independence contd.

- Write out full joint distribution using chain rule:  
$$\begin{aligned} P(\text{Toothache}, \text{Catch}, \text{Cavity}) &= P(\text{Toothache} \mid \text{Catch}, \text{Cavity}) P(\text{Catch}, \text{Cavity}) \\ &= P(\text{Toothache} \mid \text{Catch}, \text{Cavity}) P(\text{Catch} \mid \text{Cavity}) P(\text{Cavity}) \\ &= P(\text{Toothache} \mid \text{Cavity}) P(\text{Catch} \mid \text{Cavity}) P(\text{Cavity}) \end{aligned}$$
- 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.

## 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