

Uncertainty

Chapter 13

Mausam

(Based on slides by UW-AI faculty)

Knowledge Representation

KR Language	Ontological Commitment	Epistemological Commitment
Propositional Logic	facts	true, false, unknown
First Order Logic	facts, objects, relations	true, false, unknown
Temporal Logic	facts, objects, relations, times	true, false, unknown
Probability Theory	facts	degree of belief
Fuzzy Logic	facts, degree of truth	known interval values

Probabilistic Relational Models

- combine probability and first order logic

Need for Reasoning w/ Uncertainty

- The world is full of uncertainty
 - chance nodes/sensor noise/actuator error/partial info..
 - Logic is brittle
 - can't encode exceptions to rules
 - can't encode statistical properties in a domain
 - Computers need to be able to handle uncertainty
- Probability: new foundation for AI (& CS!)
- Massive amounts of data around today
 - Statistics and CS are both about data
 - Statistics lets us summarize and understand it
 - Statistics is the basis for most learning
- Statistics lets data do our work for us

Logic vs. Probability

Symbol: $Q, R \dots$	Random variable: $Q \dots$
Boolean values: T, F	Domain: you specify e.g. {heads, tails} [1, 6]
State of the world: Assignment to $Q, R \dots Z$	Atomic event: complete specification of world: $Q \dots Z$ <ul style="list-style-type: none">• Mutually exclusive• Exhaustive
	Prior probability (aka Unconditional prob: $P(Q)$)
	Joint distribution: Prob. of every atomic event

Probability Basics

- Begin with a set S : the **sample space**
 - e.g., 6 possible rolls of a die.
- $x \in S$ is a **sample point/possible world/atomic event**
- A **probability space** or **probability model** is a sample space with an assignment $P(x)$ for every x s.t.
 $0 \leq P(x) \leq 1$ and $\sum P(x) = 1$
- An **event** A is any subset of S
 - e.g. $A = \text{'die roll } < 4\text{'}$
- A **random variable** is a function from sample points to some range, e.g., the reals or Booleans

Types of Probability Spaces

Propositional or Boolean random variables

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

Discrete random variables (*finite* or *infinite*)

e.g., *Weather* is one of $\{sunny, rain, cloudy, snow\}$

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.

Arbitrary Boolean combinations of basic propositions

Axioms of Probability Theory

- All probabilities between 0 and 1
 - $0 \leq P(A) \leq 1$
 - $P(\text{true}) = 1$
 - $P(\text{false}) = 0$.
- The probability of disjunction is:

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



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:

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

Joint probability distribution for a set of r.v.s gives the probability of every atomic event on those r.v.s

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

Joint distribution can answer any question

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 there is 80% chance of cavity

- Notation for conditional distributions:

$\mathbf{P}(\text{Cavity} \mid \text{Toothache}) = 2\text{-element vector of } 2\text{-element vectors}$

- 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:


$P(\text{cavity} \mid \text{toothache}, \text{sunny}) = P(\text{cavity} \mid \text{toothache}) = 0.8$

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

Conditional Probability

- $P(A | B)$ is the probability of A given B
- Assumes that B is the only info known.
- Defined by:

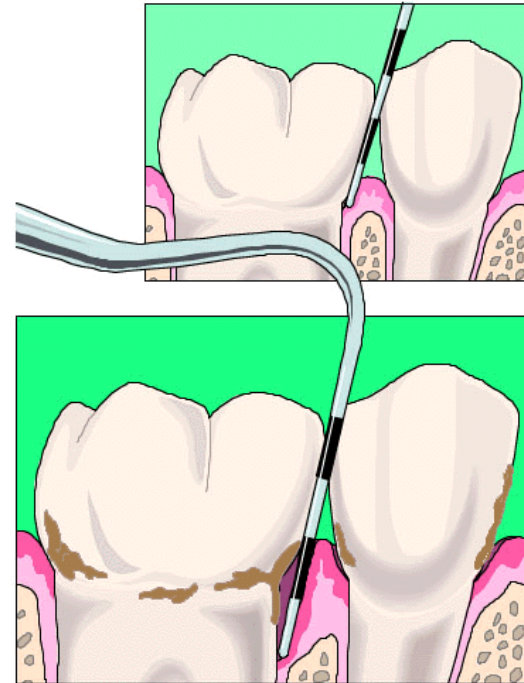
$$P(A | B) = \frac{P(A \wedge B)}{P(B)}$$



Chain Rule/Product Rule

- $$P(X_1, \dots, X_n) = P(X_n | X_1 \dots X_{n-1})P(X_{n-1} | X_1 \dots X_{n-2}) \dots P(X_1)$$
$$= \prod P(X_i | X_1, \dots, X_{i-1})$$

Dilemma at the Dentist's



What is the probability of a cavity given a toothache?
What is the probability of a cavity given the probe catches?

Inference by Enumeration

Start with the joint 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

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

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

$$\begin{aligned} P(\text{toothache}) &= .108 + .012 + .016 + .064 \\ &= .20 \text{ or } 20\% \end{aligned}$$

Inference by Enumeration

Start with the joint 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

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

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


$$P(\text{toothache} \vee \text{cavity}) = .20 + .072 + .008$$
$$.28$$

Inference by Enumeration

Start with the joint 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:

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


Complexity of Enumeration

- Worst case time: $O(d^n)$
 - Where $d = \text{max arity}$
 - And $n = \text{number of random variables}$
- Space complexity also $O(d^n)$
 - Size of joint distribution
- Prohibitive!

Independence

- A and B are *independent* iff:

$$P(A | B) = P(A)$$

$$P(B | A) = P(B)$$



These two constraints are logically equivalent

- Therefore, if A and B are independent:


$$P(A | B) = \frac{P(A \wedge B)}{P(B)} = P(A)$$

$$P(A \wedge B) = P(A)P(B)$$

Independence

A and B are independent iff

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



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

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

Complete independence is powerful but rare
What to do if it doesn't hold?

Conditional Independence

$\mathbf{P}(\textit{Toothache}, \textit{Cavity}, \textit{Catch})$ 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:

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

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

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

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

$$\mathbf{P}(\textit{Catch}|\textit{Toothache}, \textit{Cavity}) = \mathbf{P}(\textit{Catch}|\textit{Cavity})$$

Instead of 7 entries, only need 5

Conditional Independence II

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

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

Equivalent statements:

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

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

Why only 5 entries in table?

Write out full joint distribution using chain rule:

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

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

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

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

i.e., $2 + 2 + 1 = 5$ independent numbers (equations 1 and 2 remove 2)

Power of Cond. Independence

- Often, using conditional independence reduces the storage complexity of the joint distribution from exponential to linear!!
- Conditional independence is the most basic & robust form of knowledge about uncertain environments.

Bayes Rule

Bayes rules!



posterior

$$P(x, y) = P(x | y)P(y) = P(y | x)P(x)$$

$$P(x | y) = \frac{P(y | x) P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Useful for assessing **diagnostic** probability from **causal** probability:

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

Computing Diagnostic Prob. from Causal Prob.

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

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

$$P(M) = 0.0001,$$

$$P(S) = 0.1,$$

$$P(S|M) = 0.8$$

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

Note: posterior probability of meningitis still very small!

Other forms of Bayes Rule

$$P(x|y) = \frac{P(y|x) P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

$$P(x|y) = \frac{P(y|x) P(x)}{\sum_x P(y|x) P(x)}$$

$$P(x|y) = \alpha P(y|x) P(x)$$

posterior \propto likelihood \cdot prior

Conditional Bayes Rule

$$P(x | y, z) = \frac{P(y | x, z) P(x | z)}{P(y | z)}$$

$$P(x | y, z) = \frac{P(y | x, z) P(x, z)}{\sum_x P(y | x, z) P(x | z)}$$

$$P(x | y, z) = \alpha P(y | x, z) P(x | z)$$

Bayes' Rule & Cond. Independence

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

This is an example of a *naive Bayes* model:


$$\mathbf{P}(Cause, Effect_1, \dots, Effect_n) = \mathbf{P}(Cause) \prod_i \mathbf{P}(Effect_i|Cause)$$



Total number of parameters is *linear* in n

Simple Example of State Estimation

- Suppose a robot obtains measurement z
- What is $P(\text{doorOpen} | z)$?



Causal vs. Diagnostic Reasoning

- $P(open|z)$ is **diagnostic**.
- $P(z|open)$ is **causal**.
- Often **causal** knowledge is easier to obtain **count frequencies!**
- Bayes rule allows us to use causal knowledge:

$$P(open | z) = \frac{P(z | open)P(open)}{P(z)}$$

Example

- $P(z/open) = 0.6$ $P(z/\neg open) = 0.3$
- $P(open) = P(\neg open) = 0.5$

$$P(open | z) = \frac{P(z | open)P(open)}{P(z | open)p(open) + P(z | \neg open)p(\neg open)}$$

$$P(open | z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{2}{3} = 0.67$$

- z raises the probability that the door is open.

Combining Evidence

- Suppose our robot obtains another observation z_2 .
- How can we integrate this new information?
- More generally, how can we estimate $P(x / z_1 \dots z_n)$?

Recursive Bayesian Updating

$$P(x | z_1, \dots, z_n) = \frac{P(z_n | x, z_1, \dots, z_{n-1}) P(x | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})}$$

Markov assumption: z_n is independent of z_1, \dots, z_{n-1} if we know x .

$$P(x | z_1, \dots, z_n) = \frac{P(z_n | x, z_1, \dots, z_{n-1}) P(x | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})}$$

$$= \frac{P(z_n | x) P(x | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})}$$

$$= \alpha P(z_n | x) P(x | z_1, \dots, z_{n-1})$$

$$= \alpha_{1\dots n} \prod_{i=1}^n P(z_i | x) P(x)$$

Example: Second Measurement

- $P(z_2/open) = 0.5$ $P(z_2/\neg open) = 0.6$
- $P(open/z_1) = 2/3$

$$\begin{aligned} P(open | z_2, z_1) &= \frac{P(z_2 | open) P(open | z_1)}{P(z_2 | open) P(open | z_1) + P(z_2 | \neg open) P(\neg open | z_1)} \\ &= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625 \end{aligned}$$

- z_2 lowers the probability that the door is open.

These calculations seem laborious to do for each problem domain – is there a general representation scheme for probabilistic inference?



Yes - Bayesian Networks