

Limitations of First-Order Logic

- FOL is very expressive, but...consider how to translate these:
 - "*most* students graduate in 4 years"
 - $\forall x \text{ student}(x) \rightarrow \text{duration}(\text{undergrad}(x)) \leq \text{years}(4)$ (all???)
 - "*only a few* students switch majors"
 - $\exists s, m1, m2, t1, t2 \text{ student}(s) \wedge \text{major}(s, m1, t1) \wedge \text{major}(s, m2, t2) \wedge m1 \neq m2 \wedge t1 \neq t2$ (exists???)
 - "all birds can fly, *except* penguins, stuffed birds, plastic birds, birds with broken wings..."

- The problem(s) with FOL involve expressing:
 - default rules & exceptions
 - degrees of truth
 - strength of rules

Practical needs for modeling Uncertainty in KBS

- What happens if you do not know whether an antecedent is T or F?
- neither T nor F, but 'unknown' (not allowed in Boolean logic)
- FOL treats 'unasserted' facts as "could be either T or F" when determining entailment, e.g. $\{ A \wedge B \rightarrow C, A \}$ *does not entail* C
- what we often want to do is assume the most likely state (B or $\neg B$) by default
- examples:
 - what if a doctor has to make a diagnosis before white blood cell count is available?
 - or treat a patient even if history of seizures is unknown (because they are unconscious)?

- We will show how to:
 - make default assumptions that are most likely, and derive inferences from them
 - utilize prior and conditional probabilities
 - marginalize over unknowns (which is like weighted averaging by conditional probability of T or F)

Limitation of First-Order Logic

- FOL is not good at handling exceptions
 - universal quantifier means ALL; can't say "most" birds fly
 - $\forall x \text{ bird}(x) \rightarrow \text{flies}(x)$
 - asserting $\text{bird}(\text{opus}) \wedge \neg \text{flies}(\text{opus})$ in the KB would cause it to be inconsistent
 - FOL is *monotonic*: if $\alpha \models \beta$, then $\alpha \wedge \omega \models \beta$
 - adding new facts does not undo conclusions
- we could say: $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$
- but we can't enumerate all possible exceptions
 - what about a robin with a broken wing?
 - what about birds that are made out of plastic?
 - what about Big Bird?



- Uncertainty in reasoning about actions:
 - If a gun is loaded and you pull the trigger, the gun will fire, right?
 - ...unless it is a toy gun
 - ...unless it is defective
 - ...unless it is underwater
 - ...unless the barrel is filled with concrete

Possible Solutions

- Rule Strengths
- Semantic Networks
- Default Logic/Non-monotonic logics
- Closed-World Assumption and Negation-as-failure in PROLOG
- Fuzzy Logic
- Bayesian Probability

Add rule *strengths* or *priorities*

- label each rule with a number indicating its "strength" or "degree of belief"
- stronger rules override conclusions from weaker rules

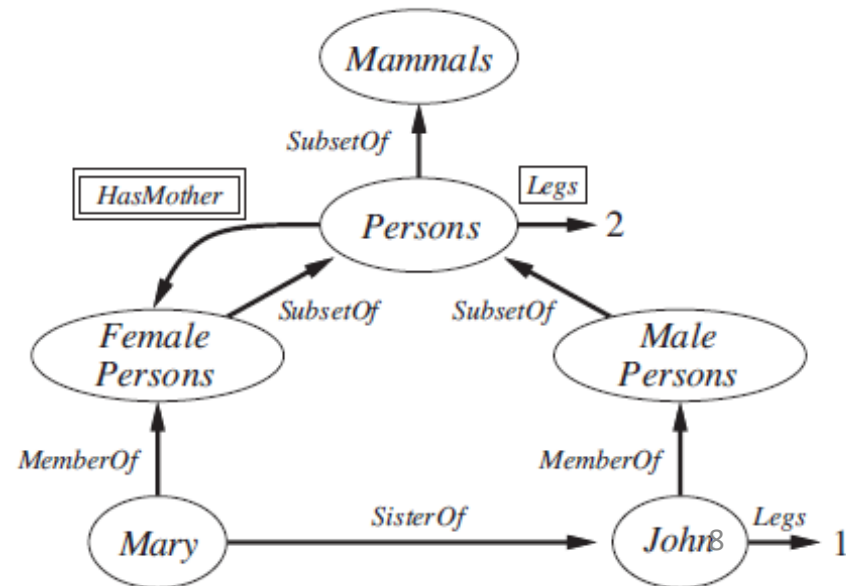
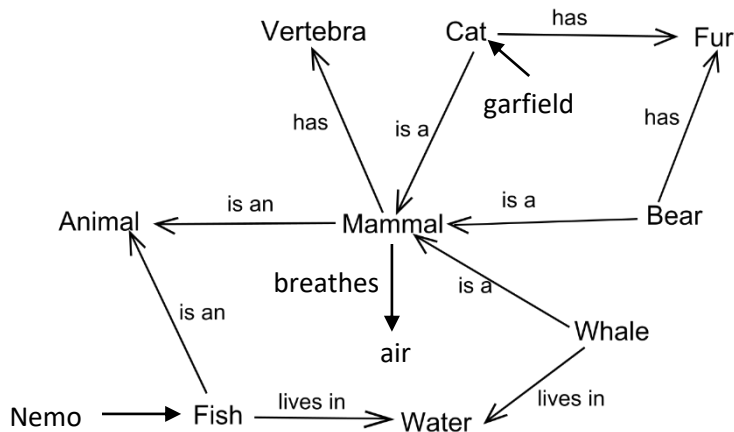
$\text{penguin}(x) \rightarrow_{0.9} \neg \text{flies}(x)$

$\text{bird}(x) \rightarrow_{0.5} \text{flies}(x)$

- an old ad-hoc approach (with unclear semantics)
- common approach in early Expert Systems
- "salience" attribute of rules in CLIPS

Semantic Networks

- graphical representation of knowledge
- nodes, slots, edges, "isa" links
- procedural mechanism for answering queries
 - follow links to get answers
 - different than formal definition of "entailment"
- inheritance
 - can override defaults



Semantic Networks

- semantic nets are a nice, graphical way of representing information
- an advantage is how they handle default info
- but there are different variations on the graphical symbology and how to express different things (like negation, universal, existence info)
 - difference between thin, thick, and dashed arcs?
 - how to express "safeDrivers are drivers who haven't been in an accident" graphically?
- what does a particular Sem Net formalism *mean*???
- try to translate it into a logic. (need more than FOL)

Non-monotonic Logics

- allow retractions later (popular for truth-maintenance systems)
- "birds fly", "penguins are birds that don't fly"
 - $\forall x \text{ bird}(x) \rightarrow \text{fly}(x)$
 - $\forall x \text{ penguin}(x) \rightarrow \text{bird}(x), \forall x \text{ penguin}(x) \rightarrow \neg \text{fly}(x)$
 - $\{\text{bird}(\text{tweety}), \text{bird}(\text{opus})\} \models \text{fly}(\text{opus})$
 - later, add that opus is a penguin, change inference
 - $\text{penguin}(\text{opus}) \models \neg \text{fly}(\text{opus})$
- Definition: A logic is *monotonic* if everything that is entailed by a set of sentences α is entailed by any superset of sentences $\alpha \wedge \beta$
 - opus example is *non-monotonic*


Default Logic

- example syntax of a default rule:
 - $\text{bird}(x): \text{fly}(x) / \text{fly}(x)$ or $\text{bird}(x) \succ \text{fly}(x)$
 - analogous to $\forall x \text{bird}(x) \rightarrow \text{fly}(x)$, but allows exceptions
 - meaning: "if PREMISE is satisfied and it is not inconsistent to believe CONSEQUENT, then CONSEQUENT"
 - $\{\text{bird}(\text{tweet}), \text{bird}(\text{opus}), \neg \text{fly}(\text{opus}), \text{bird}(x): \text{fly}(x) / \text{fly}(x)\}$
 $\models \{\text{fly}(\text{tweet}), \neg \text{fly}(\text{opus})\}$
- requires *fixed-point semantics* (different model theory and inference procedures)

Circumscription

- an alternative approach to default logic
- add *abnormal* predicates to rules
 - $\forall x \text{ bird}(x) \wedge \neg \text{abnormal}_1(x) \rightarrow \text{fly}(x)$
 - $\forall x \text{ penguin}(x) \wedge \neg \text{abnormal}_2(x) \rightarrow \text{bird}(x)$
 - $\forall x \text{ penguin}(x) \wedge \neg \text{abnormal}_3(x) \rightarrow \neg \text{fly}(x)$
- algorithm: minimize the number of *abnormals* needed to make the KB consistent
 - $\{\text{bird}(\text{tweety}), \text{fly}(\text{tweety}), \text{bird}(\text{opus}), \text{penguin}(\text{opus}), \neg \text{fly}(\text{opus})\}$ is INCONSISTENT
 - $\{\text{bird}(\text{tweety}), \text{fly}(\text{tweety}), \text{bird}(\text{opus}), \text{penguin}(\text{opus}), \neg \text{fly}(\text{opus}), \text{abnormal}_1(\text{opus})\}$ is CONSISTENT


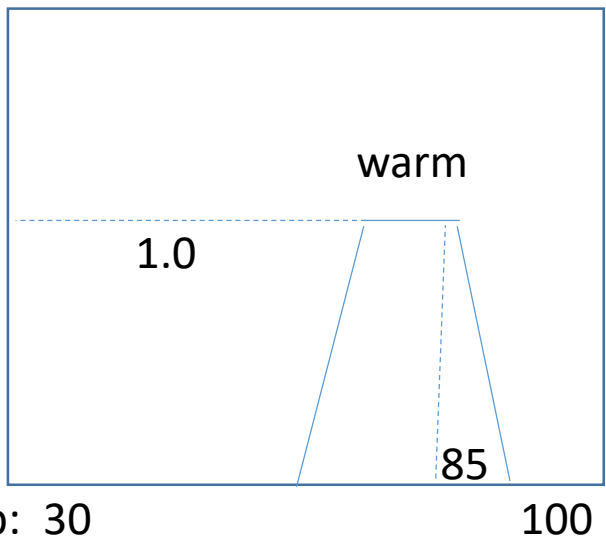
Fuzzy Logic




- some expressions involve "degrees" of truth, like "John is tall"
- membership functions
- "most students with high SATs have high GPAs"
- inference by computing with membership funcs.
 - "only days that are warm and not windy are good for playing frisbee"
 - suppose today is 85 and the wind is 15 kts NE
 - $T(A \wedge B) = \min(T(A), T(B))$
 - $T(A \vee B) = \max(T(A), T(B))$
- popular for control applications (like thermostats...)

Fuzzy Logic

- doing inference in FL involve computing truncation (min) and intersection with membership functions
 - i.e. to evaluate satisfaction of antecedents of a rule
 - (temp is warm) and (wind is not-windy) -> playFrisbee



2D: tempXwind



Handling Defaults in Prolog

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$
 - `bird(tweety)`
 - `bird(woodstock)`
 - `bird(opus)` `penguin(opus)`

initial KB has 4 facts

Handling Defaults in Prolog

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$

- $\text{bird}(\text{tweety})$
- $\text{bird}(\text{woodstock})$
- $\text{bird}(\text{opus})$ $\text{penguin}(\text{opus})$

initial KB has 4 facts

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$

- $\text{bird}(\text{tweety})$ $\neg \text{penguin}(\text{tweety})$
- $\text{bird}(\text{woodstock})$ $\neg \text{penguin}(\text{woodstock})$
- $\text{bird}(\text{opus}),$ $\text{penguin}(\text{opus})$

the problem here is that, if you add a qualifying condition like $\neg \text{penguin}$ to a rule in FOL, then you have to explicitly say whether every individual is a penguin or not (which is not scalable to large KBs)

Handling Defaults in Prolog

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$

- $\text{bird}(\text{tweety})$
- $\text{bird}(\text{woodstock})$
- $\text{bird}(\text{opus})$ $\text{penguin}(\text{opus})$

initial KB has 4 facts

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \rightarrow \text{flies}(x)$

- $\text{bird}(\text{tweety})$ $\neg \text{penguin}(\text{tweety})$
- $\text{bird}(\text{woodstock})$ $\neg \text{penguin}(\text{woodstock})$
- $\text{bird}(\text{opus}),$ $\text{penguin}(\text{opus})$

the problem here is that, if you add a qualifying condition like $\neg \text{penguin}$ to a rule in FOL, then you have to explicitly say whether every individual is a penguin or not

- $\forall x \text{ bird}(x) \wedge \neg \text{penguin}(x) \wedge \neg \text{emu}(x) \rightarrow \text{flies}(x)$

- $\text{bird}(\text{tweety})$ $\neg \text{penguin}(\text{tweety})$ $\neg \text{emu}(\text{tweety})$
- $\text{bird}(\text{woodstock})$ $\neg \text{penguin}(\text{woodstock})$ $\neg \text{emu}(\text{woodstock})$
- $\text{bird}(\text{opus}),$ $\text{penguin}(\text{opus})$ $\neg \text{emu}(\text{opus})$

which is not scalable to large KBs

if we add another condition like $\neg \text{emu}$, then we have explicitly identify all the non-emus

- other examples:

- a football player is eligible to play in a game, *unless* they have not passed a physical, or are on academic probation
- an item is on sale (50% off), *unless* it is already discounted
- a house can be sold, as long as it does not have a lien on it
- fish is a healthy option for protein, *unless* it has high mercury levels (shark, swordfish, orange roughy...)
- in all these cases, you would have to add a negative antecedent to a FOL rule, but then have to assert things like \neg academic_probation(<player>) for all players, or \neg highMg(trout), \neg highMg(bass), \neg highMg(catfish)...

$A \wedge B \rightarrow C \dots \neg A \vee \neg B \vee C$

$A \wedge \neg B \rightarrow C \dots \neg A \vee B \vee C$

Handling Defaults in Prolog

- Potential problems:
 - 1) can't assert negative facts, e.g. `¬penguin(tweety)`
 - 2) can't have negative literals as antecedents in definite clauses

```
dog(fido) .
```

```
dog(snoopy) .
```

```
canary(tweety) .
```

```
canary(woodstock) .
```

```
penguin(opus) .
```

```
animal(X) :- mammal(X) .
```

```
animal(X) :- bird(X) .
```

```
dangerous_animal(X) :- animal(X), has_sharp_teeth(X), aggressive(X) .
```

Closed-World Assumption (CWA) in PROLOG

- every fact that is not explicitly asserted (or provable) is *assumed to be false*
- can include negated antecedents in rules (" $\backslash+$ " = not)

```
stench(1,2) . // col,row
```

```
stench(2,3) .
```

```
stench(1,4) .
```

```
wumpus_free(X,Y) :- room(X,Y), adjacent(X,Y,P,Q), \+ stench(P,Q) .
```

```
?- wumpus_free(1,3) .
```

No

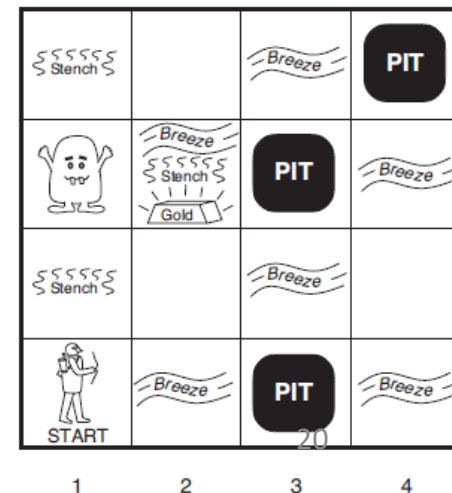
```
?- wumpus_free(2,2) .
```

Yes (because no stench in 2,1)

this prolog rule is equivalent to:

$\forall x,y,p,q \text{ room}(x,y) \wedge \text{adjacent}(x,y,p,q) \wedge \neg \text{stench}(p,q) \rightarrow \text{wumpus_free}(x,y)$

which is not a definite clause (because CNF has 2 positive literals), so technically, we could not do back-chaining; can't put $\neg \text{stench}$ on goal stack



- using CWA for default reasoning

```
bird(X) :- canary(X).
```

```
?- bird(tweety). Yes
```

```
bird(X) :- penguin(X).
```

```
?- canFly(tweety). Yes
```

```
canary(tweety).
```

```
?- bird(opus). Yes
```

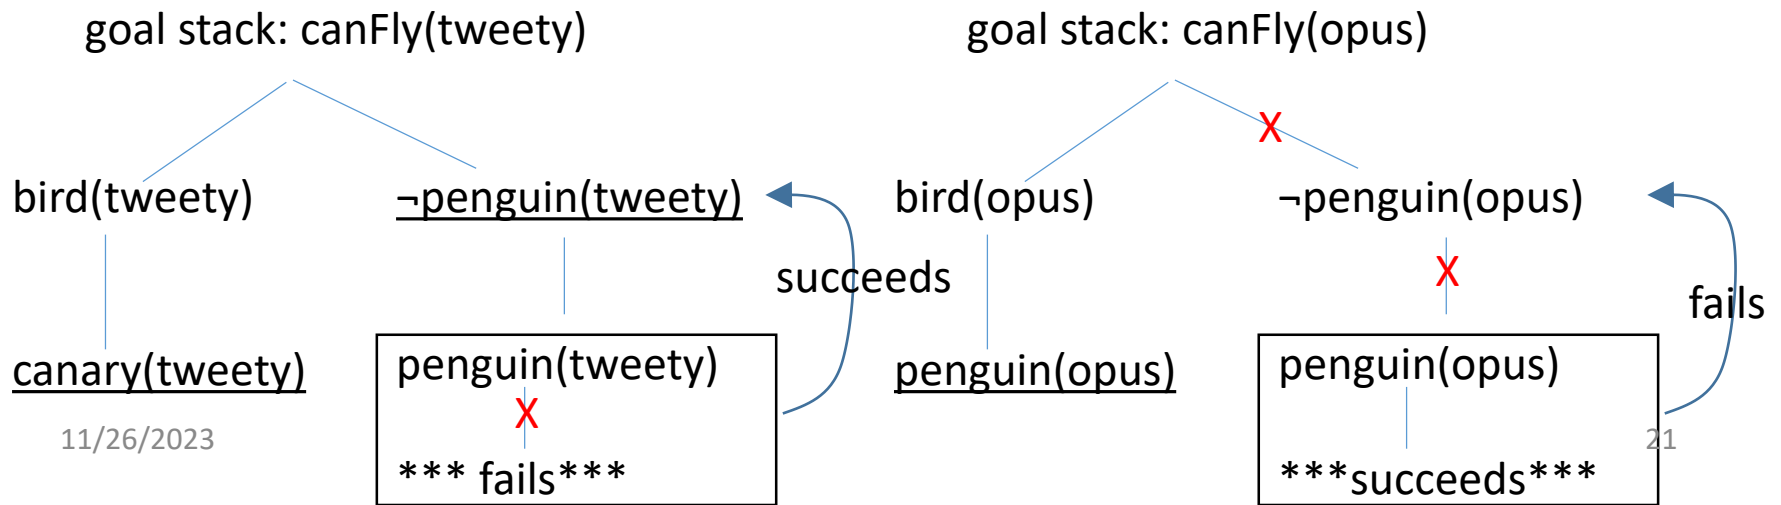
```
penguin(opus).
```

```
?- canFly(opus). No
```

```
canFly(X) :- bird(X), \+ penguin(X).
```

- how is negation-as-failure implemented?

- modify back-chaining to handle negative antecedents
- when trying to prove $\neg P(X)$ on goal stack, try proving $P(X)$ and if fail then $\neg P(X)$ succeeds



Probability (Ch. 12)

- an alternative route to encoding default rules like "most birds fly" is to quantify it using probability, $p(\text{fly} | \text{bird})=0.95$
- probabilistic reasoning has had a major impact on AI over the years
 - conferences and journals on UAI (Uncertainty in AI)
- probabilistic models has led to major algorithms like:
 - Hidden Markov Models (applications to speech, genomics...)
 - SLAM (simultaneous localization and mapping) for robotics
 - Bayesian networks/graphical models (as knowledge bases)
 - Kalman filters, ICA, POMPDs, ...
 - Reinforcement Learning

Axioms of Probability

- for event e : $0 \leq P(e) \leq 1$
- for mutually exclusive events $e_1 \dots e_n$: $\sum_i P(e_i) = 1$
- negation: $P(\neg e) = 1 - P(e)$
- Kolmogorov axiom for non-exclusive events:
$$P(a \vee b) = P(a) + P(b) - P(a, b)$$

Prior and Conditional Probabilities

- encode knowledge in the form of *prior* probabilities and *conditional* probabilities
 - $P(x \text{ speaks portugese})=0.012$
 - $P(x \text{ is from Brazil})=0.007$
 - $P(x \text{ speaks portugese} | x \text{ is from Brazil})=0.9$
 - $P(x \text{ flies} | x \text{ is a bird})=0.9$ (?)
- inference is done by calculating *posterior* probabilities given evidence (using Bayes' Rule)
 - compute $P(\text{cavity} | \text{toothache, flossing, dental history, recent consumption of candy...})$
 - compute $P(\text{fed will raise interest rate} | \text{unemployment}=5\%, \text{inflation}=0.5\%, \text{GDP}=2\%, \text{recent geopolitical events...})$

prior probs

conditional probs

Bayes' Rule

- product rule : *joint prob* $P(A,B) = P(A | B) * P(B)$
 - $P(A | B)$ is read as "probability of A *given* B"
 - in general, $P(A,B) \neq P(A) * P(B)$ (unless A and B are independent)
- Bayes' Rule: convert between causal and diagnostic

$$P(H | E) = \frac{P(E | H) \cdot P(H)}{P(E)}$$

H = hypothesis (cause, disease)
E = evidence (effect, symptoms)

- joint probabilities: $P(E,H)$, priors: $P(H)$
- conditional probabilities play role of "rules"
 - people with a toothache are likely to have a cavity
 - $p(\text{cavity} | \text{toothache}) = 0.6$

Causal vs. diagnostic knowledge

- *causal*: $P(x \text{ has a toothache} | x \text{ has a cavity})=0.9$
- *diagnostic*: $P(x \text{ has a cavity} | x \text{ has a toothache})=0.6$
- typically it is easier to articulate knowledge in the causal direction, but we often want to use it in a diagnostic way to make inferences from observations

- Joint probability table (JPT)
 - you can calculate answer to any question from JPT
 - the problem is there are exponential # of entries (2^N , where N is the number of binary random variables)

	<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

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

- Joint probability table (JPT)
 - you can calculate answer to any question from JPT
 - the problem is there are exponential # of entries (2^N , where N is the number of binary random variables)

	<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

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

- Joint probability table (JPT)
 - you can calculate answer to any question from JPT
 - the problem is there are exponential # of entries (2^N , where N is the number of binary random variables)

	<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

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

- **marginalization** - summing out unknown variables

$$P(cavity) = P(cavity, toothache, catch) + P(cavity, toothache, \neg catch) + P(cavity, \neg toothache, catch) + P(cavity, \neg toothache, \neg catch)$$

$$P(cavity) = 0.108 + 0.012 + 0.072 + 0.008 = 0.2$$

	<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

=0.2

- **normalization**

- suppose we want to compute a *conditional* prob, like $P(X|Y,Z)$
- using the product rule, we could calculate it using *joint* probs:
 - $P(X|Y,Z) = P(X,Y,Z)/P(Y,Z)$
- would have to marginalize over X to compute the denominator
 - $P(Y,Z) = P(X,Y,Z)+P(-X,Y,Z)$
- a simpler way to calculate the conditional prob is to compute 2 joint probabilities, $P(X,Y,Z)$ and $P(-X,Y,Z)$, and normalize them so they sum up to 1 (X has to be T or F in context of Y and Z)
- this represents the evidence "for" and "against" X, given Y and Z
 - $P(X|Y,Z) = \alpha P(X,Y,Z) ; \alpha=1/(P(X,Y,Z)+P(-X,Y,Z))$
- since we have to compute probs both *for* and *against*, it is conventional to represent them as a vector:
 - $\langle P(X,Y,Z), P(-X,Y,Z) \rangle$
- technically, they don't add up to 1, but we can make them sum to one by dividing by the sum to normalize them
 - $\alpha \langle P(X,Y,Z), P(-X,Y,Z) \rangle ; \alpha=1/(P(X,Y,Z)+P(-X,Y,Z))$
 - $P(X|Y,Z) = P(X,Y,Z)/(P(X,Y,Z)+P(-X,Y,Z))$

Conditional Independence

- Applying Bayes' Rule in larger domains has a scalability problem
 - the size of the JPT grows exponentially with the number of variables (2^n for n variables)
- Solution to reduce complexity:
 - employ the Independence Assumption
- Most variables are not strictly independent; most variables are at least partially correlated (but which is cause and which is effect?).
- However, many variables are *conditionally* independent.

A and B are *conditionally independent given C* if:

$$P(A, B | C) = P(A | C)P(B | C), \text{ or equivalently}$$

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

Conditional Independence

If I have a cavity, the probability that the probe catches in it doesn't depend on whether I have a toothache:

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

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*:

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

- conditional independence gives us an efficient way to combine evidence
 - consider $P(\text{Cav} | \text{toothache}, \text{catch})$
 - using Bayes' Rule:
 - $P(\text{Cav} | \text{toothache}, \text{catch}) \propto P(\text{toothache} \wedge \text{catch} | \text{Cav}) * P(\text{Cav})$
 - this requires a mini JPT for all combinations of evidence
 - assuming toothache is conditionally independent of catch given Cavity:
 - $P(\text{toothache} \wedge \text{catch} | \text{Cav}) = P(\text{toothache} | \text{Cav}) * P(\text{catch} | \text{Cav})$
 - therefore...

$$P(\text{Cav} | \text{toothache}, \text{catch}) \propto P(\text{toothache} | \text{Cav}) * P(\text{catch} | \text{Cav}) * P(\text{Cav})$$

Naive Bayes algorithm

- suppose you have a phenomenon that causes several different effects that could be observed
- Cause \rightarrow Effect₁, Effect₂, ..., Effect_n
- each effect is probabilistic, but *assume they are all conditionally independent of each other*
- Then an efficient method for detecting or classifying probable causes is:

$$P(\text{Cause}, \text{Effect}_1, \dots, \text{Effect}_n) = P(\text{Cause}) \prod_i P(\text{Effect}_i | \text{Cause})$$

- if you have some unobserved vars (y), could marginalize them out, but it leads to same Eqn above

$$P(\text{Cause} | \mathbf{e}) = \alpha \sum_y P(\text{Cause}, \mathbf{e}, y)$$

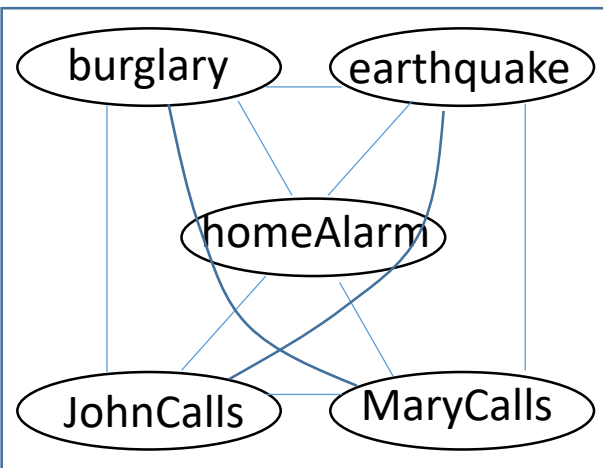
- Example: classifying documents as Bag-of-Words

- $P(\text{doctype}=\text{sports} | \text{words}) = P(\text{sports}) * (\text{has "score"} | \text{sports}) * (\text{has "referee"} | \text{sports}) * \dots$


Bayesian Networks (Sec. 13.1 and the first page of Sec 13.2)

- graphical models where *edges represent conditional probabilities*
 - efficient representation because missing edges are assumed to be *conditionally independent* given the nodes in between
- popular for modern AI systems (expert systems)
 - important for handling uncertainty

all vars are correlated, $O(n^2)$ edges, requires full JPT with 2^n rows




Naive Bayes: compute probability of 1 var depending on all the others (n-1)



requires *independence assumption*

Bayesian Network: selected edges represent conditional dependence




more natural: links follow causality

Bayesian Networks (Sec. 13.1-2)

- prob of each node depends on parents; specify with a mini-JPT
- full JPT has $2^5=32$ entries - can answer any query from JPT
- joint prob of full state $\langle j, m, a, \neg b, \neg e \rangle$ is product of prob over all nodes like
- prob of each node is conditioned on parents

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i | \text{parents}(X_i))$$




$$\begin{aligned}
 &P(j \wedge m \wedge a \wedge \neg b \wedge \neg e) \\
 &= P(j|a)P(m|a)P(a|\neg b, \neg e)P(\neg b)P(\neg e) \\
 &= 0.9 \times 0.7 \times 0.001 \times 0.999 \times 0.998 \\
 &\approx 0.00063
 \end{aligned}$$

- Efficient algorithms for computing inferences or outcomes conditioned on observations/evidence
 - Variable elimination: factor computations into a tree of products and sums (algebraic calculation from formula)
 - rearrange to minimize number of adds and mults...

$P(\text{Burglary} | \text{JohnCalls} = \text{true}, \text{MaryCalls} = \text{true})$

$$P(b | j, m) = \alpha \sum_e \sum_a P(b)P(e)P(a|b,e)P(j|a)P(m|a)$$

$$P(b | j, m) = \alpha P(b) \sum_e P(e) \sum_a P(a|b,e)P(j|a)P(m|a)$$



- Belief propagation: graph algorithm that updates probs of neighboring nodes when belief of any node changes





Figure 13.9 A Bayesian network for evaluating car insurance applications.

Initial evidence: car won't start

Testable variables (green), "broken, so fix it" variables (orange)

Hidden variables (gray) ensure sparse structure, reduce parameters



- Many modern knowledge-based systems are based on probabilistic inference
 - including Bayesian networks, Hidden Markov Models, (HMMs), Markov Decision Problems (MDPs)
 - example: Bayesian networks are used for inferring user goals or help needs from actions like mouse clicks in an automated software help system (think 'Clippy')
 - *Decision Theory* combines *utilities* with *probabilities* of outcomes to decide actions to take
- the challenge is capturing all the numbers needed for the prior and conditional probabilities
 - objectivists (frequentists) - probabilities represent outcomes of trials/experiments
 - subjectivists - probabilities are degrees of belief
- probability and statistics is at the core of many Machine Learning algorithms

