Probabilistic representation

Applied artificial intelligence (EDA132)
Lecture "10"
2012-04-19
Elin A. Topp

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Uncertainty

Situation: Get to the airport in time for the flight (by car)

Action A_t := "Leave for airport t minutes before flight departs"

Question: will A_t get me there on time?

Deal with:

1) partial observability (road states, other drivers, ...)

2) noisy sensors (traffic reports)

3) uncertainty in action outcomes (flat tire, car failure, ...)

4) complexity of modeling and predicting traffic

Use pure logic? Well... :

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

or 2) leads to conclusions too weak for decision making:

"A_{25} will get me there on time if there is no accident and it does not rain and my tires hold, and ..."

(A_{1440} would probably hold, but the waiting time would be intolerable, given the quality of airport food...)

Rational decision

A_{25}, A_{90}, A_{180}, A_{1440}, ... what is "the right thing to do?"

Obviously dependent on relative importance of goals (being in time vs minimizing waiting time) AND on their respective likelihood of being achieved.

Uncertain reasoning: diagnosing a patient, i.e., find the CAUSE for the symptoms displayed.

"Diagnostic" rule: Toothache ⇒ Cavity ???

Complex rule: Toothache ⇒ Cavity ∨ GumProblem ∨ Abscess ∨ ... ???

"Causal" rule: Cavity ⇒ Toothache ???

Rational decision

A_{25}, A_{90}, A_{180}, A_{1440}, ... what is "the right thing to do?"

Obviously dependent on relative importance of goals (being in time vs minimizing waiting time) AND on their respective likelihood of being achieved.

Uncertain reasoning: diagnosing a patient, i.e., find the CAUSE for the symptoms displayed.

"Diagnostic" rule: Toothache ⇒ Cavity ??? No!

Complex rule: Toothache ⇒ Cavity ∨ GumProblem ∨ Abscess ∨ ... ???

"Causal" rule: Cavity ⇒ Toothache ???

Rational decision

A_{25}, A_{90}, A_{180}, A_{1440}, ... what is "the right thing to do?"

Obviously dependent on relative importance of goals (being in time vs minimizing waiting time) AND on their respective likelihood of being achieved.

Uncertain reasoning: diagnosing a patient, i.e., find the CAUSE for the symptoms displayed.

"Diagnostic" rule: Toothache ⇒ Cavity ??? No!

Complex rule: Toothache ⇒ Cavity ∨ GumProblem ∨ Abscess ∨ ... ??? Too much!

"Causal" rule: Cavity ⇒ Toothache ???

Rational decision

A_{25}, A_{90}, A_{180}, A_{1440}, ... what is "the right thing to do?"

Obviously dependent on relative importance of goals (being in time vs minimizing waiting time) AND on their respective likelihood of being achieved.

Uncertain reasoning: diagnosing a patient, i.e., find the CAUSE for the symptoms displayed.

"Diagnostic" rule: Toothache ⇒ Cavity ??? No!

Complex rule: Toothache ⇒ Cavity ∨ GumProblem ∨ Abscess ∨ ... ??? Too much!

"Causal" rule: Cavity ⇒ Toothache ??? Well... not always

Using logic?

Fixing such "rules" would mean to make them logically exhaustive, but that is bound to fail due to:

Laziness (too much work to list all options)

Theoretical ignorance (there is simply no complete theory)

Practical ignorance (might be impossible to test exhaustively)

⇒ better use ***probabilities*** to represent certain ***knowledge states***

⇒ Rational decisions (decision theory) combine probability and utility theory

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Probability

Probabilistic assertions summarise 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 state of knowledge

e.g., $P(A_{25} \mid \text{no reported accidents}) = 0.06$

Not claims of a "probabilistic tendency" in the current situation, but maybe learned from past experience of similar situations.

Probabilities of propositions change with new evidence:

e.g., $P(A_{25} \mid \text{no reported accidents, it's 5:00 in the morning}) = 0.15$

Making decisions under uncertainty

Suppose the following believes (from past experience):

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

Which action to choose?

Depends on my preferences for "missing flight" vs. "waiting (with airport cuisine)", etc.

Utility theory is used to represent and infer preferences

Decision theory = utility theory + probability theory

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Probability basics

A set Ω - the sample space, e.g., the 6 possible rolls of a die.

$\omega \in \Omega$ is a sample point / possible world / atomic event

A probability space of probability model is a sample space with an assignment $P(\omega)$ for every $\omega \in \Omega$ so that:

$$0 \leq P(\omega) \leq 1$$

$$\sum_{\omega} P(\omega) = 1$$

An event A is any subset of Ω

$$P(A) = \sum_{\{\omega \in A\}} P(\omega)$$

E.g., $P(\text{die roll} < 4) = P(1) + P(2) + P(3) = 1/6 + 1/6 + 1/6 = 1/2$

Random variables

A random variable is a function from sample points to some range, e.g., the reals or Booleans,

e.g., *Odd(1) = true.*

P induces a *probability distribution* for any random variable *X*

$$P(X = x_i) = \sum_{\{\omega : X(\omega) = x_i\}} P(\omega)$$

e.g., $P(Odd = true) = P(1) + P(3) + P(5) = 1/6 + 1/6 + 1/6 = 1/2$

Propositions

A proposition describes the event (set of sample points) where it (the proposition) holds, i.e.,

Given Boolean random variables A and B:

event a = *set of sample points where* $A(\omega) = true$

event $\neg a$ = *set of sample points where* $A(\omega) = false$

event $a \wedge b$ = *points where* $A(\omega) = true$ *and* $B(\omega) = true$

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.

Prior probability

Prior or *unconditional probabilities* of propositions

e.g., *P(Cavity = true) = 0.2* and

P(Weather = sunny) = 0.72

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

Probability distribution gives values for all possible assignments (normalised):

P*(Weather) = ⟨0.72, 0.1, 0.08, 0.1⟩*

Joint probability distribution for a set of (independent) random variables gives the probability of every atomic event on those random variables (i.e., every sample point):

P*(Weather, Cavity)* = a *4 x 2* matrix of values:

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

Posterior probability

Most often, there is *some* information, i.e., *evidence*, that one can base their belief on:

e.g., *P(cavity) = 0.2* (prior, no evidence for anything), but

P(cavity | toothache) = 0.6

corresponds to belief *after the arrival of some evidence*
*(*also: *posterior* or *conditional probability).*

OBS: NOT "if toothache, then 60% chance of cavity"

THINK *"given that toothache is all I know" instead!*

Posterior probability

Most often, there is *some* information, i.e., *evidence*, that one can base their belief on:

e.g., *P(cavity) = 0.2* (prior, no evidence for anything), but

P(cavity | toothache) = 0.6

corresponds to belief *after the arrival of some evidence*
*(*also: *posterior* or *conditional probability).*

OBS: NOT "if toothache, then 60% chance of cavity"

THINK *"given that toothache is all I know" instead!*

Evidence remains valid after more evidence arrives, but it might become less useful

Evidence may be completely useless, i.e., irrelevant.

P(cavity | toothache, sunny) = P(cavity | toothache)

Domain knowledge lets us do this kind of inference.

Posterior probability (2)

Definition of conditional / posterior probability:

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

or as *Product rule* (for a and b being true, we need b true and then a true, given b):

$$P(a \wedge b) = P(a \mid b)\, P(b) = P(b \mid a)\, P(a)$$

and in general for whole distributions (e.g.):

$$\mathbf{P}(Weather, Cavity) = \mathbf{P}(Weather \mid Cavity)\, \mathbf{P}(Cavity)$$

(gives a *4x2* set of equations)

Chain rule (successive application of product rule):

$$\mathbf{P}(X_1, \ldots, X_n) = \mathbf{P}(X_1, \ldots, X_{n-1})\, \mathbf{P}(X_n \mid X_1, \ldots, X_{n-1})$$

$$= \mathbf{P}(X_1, \ldots, X_{n-2})\, \mathbf{P}(X_{n-1} \mid X_1, \ldots, X_{n-1})\, \mathbf{P}(X_n \mid X_1, \ldots, X_{n-1})$$

$$= \ldots = \prod_{i=1}^{n} \mathbf{P}(X_i \mid X_1, \ldots, X_{i-1})$$

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

The suicidal student

A young student kills herself. Her diary is found. In the diary she speculates about her childhood and the possibility of her father abusing her during childhood. She had reported headaches to her friends and therapist, and started the diary due to the therapist's recommendation.

The father ends up in court, since

"headaches are caused by PTSD, and PTSD is caused by abuse"

What went wrong here?

The suicidal student

A young student kills herself. Her diary is found. In the diary she speculates about her childhood and the possibility of her father abusing her during childhood. She had reported headaches to her friends and therapist, and started the diary due to the therapist's recommendation.

The father ends up in court, since

"headaches are caused by PTSD, and PTSD is caused by abuse"

What went wrong here?

Psychologist knowing the math argues:

P(headache | PTSD) = high (statistics)

P(PTSD | abuse in childhood) = high (statistics)

but:

You do not know anything (in this case) of

P(PTSD | headache)

P(abuse in childhood | headache)

with only the evidence of headache and a speculation!

Inference

Probabilistic inference:

Computation of posterior probabilities given observed evidence

starting out with the full joint distribution as "knowledge base":

Inference by enumeration

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

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

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

Inference

Probabilistic inference:

Computation of posterior probabilities given observed evidence

starting out with the full joint distribution as "knowledge base":

Inference by enumeration

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

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

Probabilistic inference:

Computation of posterior probabilities given observed evidence

starting out with the full joint distribution as "knowledge base":

Inference by enumeration

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

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

Probabilistic inference:

Computation of posterior probabilities given observed evidence

starting out with the full joint distribution as "knowledge base":

Inference by enumeration

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

Can also compute posterior probabilities:

$$P(\neg cavity \mid toothache) = \frac{P(\neg cavity \wedge toothache)}{P(toothache)}$$

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

Normalisation

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

Denominator can be viewed as a *normalisation constant*:

$$\boldsymbol{P}(\mathit{Cavity} \mid \mathit{toothache}) = \alpha\, \boldsymbol{P}(\mathit{Cavity}, \mathit{toothache})$$

$$= \alpha[\boldsymbol{P}(\mathit{Cavity}, \mathit{toothache}, \mathit{catch}) + \boldsymbol{P}(\mathit{Cavity}, \mathit{toothache}, \neg\mathit{catch})]$$

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

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

Normalisation

	toothache		¬ toothache	
	catch	¬ catch	catch	¬ catch
cavity	0.108	0.012	0.072	0.008
¬ cavity	0.016	0.064	0.144	0.576

Denominator can be viewed as a *normalisation constant*:

$$\mathbf{P}(Cavity \mid toothache) = \alpha\, \mathbf{P}(Cavity, toothache)$$

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

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

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

And the good news:

We can compute $\mathbf{P}(Cavity \mid toothache)$ without knowing the value of $P(toothache)$!

... but

n Boolean variables give us an input table of size $O(2^n)$ ...

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Independence

A and *B* are *independent* iff

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

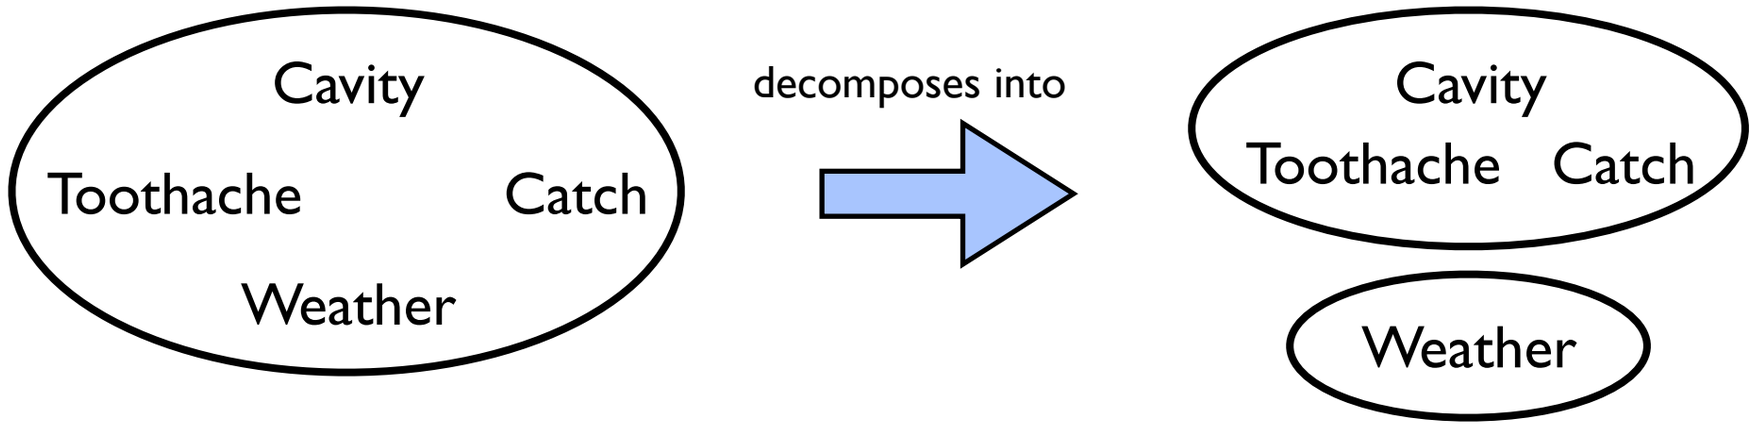

$$P(Toothache, Catch, Cavity, Weather) = P(Toothache, Catch, Cavity)\,P(Weather)$$

32 entries reduced to 8 + 4. This absolute independence is powerful but rare!

Some fields (like dentistry) have still a lot, maybe hundreds, of variables, none of them being independent.

What can be done to overcome this mess...?

Conditional independence

P*(Toothache, Cavity, Catch)* has $2^3 - 1 = 7$ independent entries (must sum up to 1)

But: If there is a cavity, the probability for "catch" does not depend on whether there is a toothache:

(1) *P(catch | toothache, cavity) = P(catch | cavity)*

The same holds when there is no cavity:

(2) *P(catch | toothache, ¬cavity) = P(catch | ¬cavity)*

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

P*(Catch | Toothache, Cavity) =* ***P****(Catch | Cavity)*

Writing out full joint distribution using chain rule:

P*(Toothache, Catch, Cavity)*
= ***P****(Toothache | Catch, Cavity)* ***P****(Catch, Cavity)*
= ***P****(Toothache | Catch, Cavity)* ***P****(Catch | Cavity)* ***P****(Cavity)*
= ***P****(Toothache | Cavity)* ***P****(Catch | Cavity)* ***P****(Cavity)*

gives thus $2 + 2 + 1 = 5$ independent entries

Conditional independence (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*.

Hence:

Conditional independence is our most basic and robust form of knowledge about uncertain environments

Bayes' Rule

Recap *product rule*: $P(a \wedge b) = P(a \mid b)\, P(b) = P(b \mid a)\, P(a)$

$$\Rightarrow \text{ Bayes' Rule } P(a \mid b) = \frac{P(b \mid a)\, P(a)}{P(b)}$$

or in distribution form:

$$\boldsymbol{P}(Y \mid X) = \frac{\boldsymbol{P}(X \mid Y)\, \boldsymbol{P}(Y)}{\boldsymbol{P}(X)} = \alpha\, \boldsymbol{P}(X \mid Y)\, \boldsymbol{P}(Y)$$

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

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

E.g., with M "meningitis", S "stiff neck":

$$P(m \mid s) = \frac{P(s \mid m)\, P(m)}{P(s)} = \frac{0.8 * 0.0001}{0.1} = 0.0008$$ (not too bad, really!)

Bayes' Rule and conditional independence

$$\boldsymbol{P}(Cavity \mid toothache \wedge catch)$$
$$= \alpha \; \boldsymbol{P}(toothache \wedge catch \mid Cavity)\, \boldsymbol{P}(Cavity)$$
$$= \alpha \; \boldsymbol{P}(toothache \mid Cavity)\, \boldsymbol{P}(catch \mid Cavity)\, \boldsymbol{P}(Cavity)$$

An example of a *naive Bayes* model:

$$\boldsymbol{P}(Cause, Effect_1, \ldots, Effect_n) = \boldsymbol{P}(Cause) \prod_i \boldsymbol{P}(Effect_i \mid Cause)$$

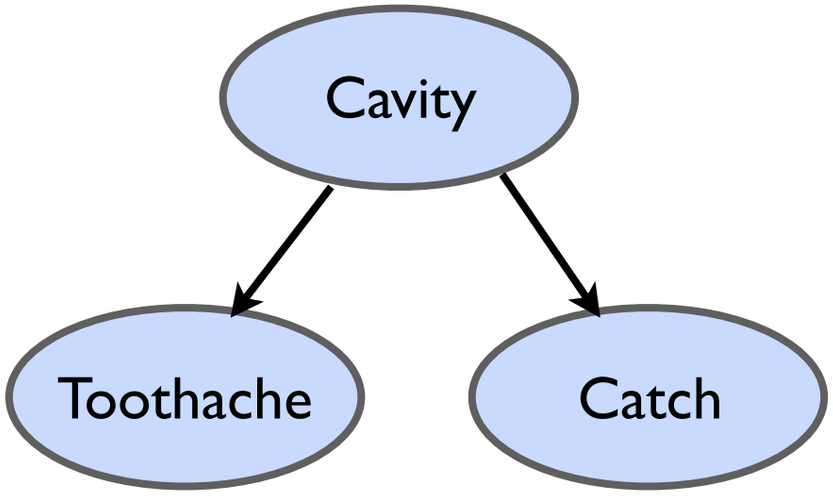

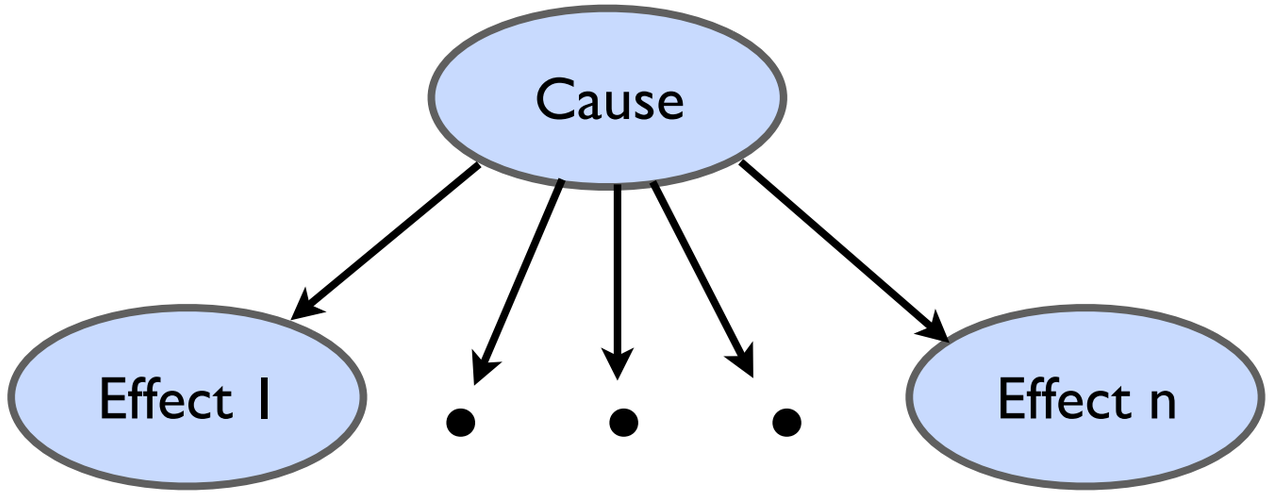

The total number of parameters is *linear* in *n*

Wumpus World

1,4	2,4	3,4	4,4
1,3	2,3	3,3	4,3
1,2 B ok	2,2	3,2	4,2
1,1 ok	2,1 B ok	3,1	4,1

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 $\boldsymbol{P}(P_{1,1}, ..., P_{4,4}, B_{1,1}, B_{1,2}, B_{2,1})$

Apply product rule: $\boldsymbol{P}(B_{1,1}, B_{1,2}, B_{2,1} \mid P_{1,1}, ..., P_{4,4},) \boldsymbol{P}(P_{1,1}, ..., P_{4,4})$

(getting $P(\mathit{Effect} \mid \mathit{Cause})$.)

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

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

$$\boldsymbol{P}(P_{1,1}, ..., P_{4,4}) = \prod_{i,j=1,1}^{4,4} \boldsymbol{P}(P_{i,j}) = 0.2^n * 0.8^{16-n}$$

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 $\mathbf{P}(P_{1,3} \mid known, b)$

Define: $Unknown = P_{i,j}$s other than $P_{1,3}$ and $Known$

For inference by enumeration, we have

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

Grows exponentially with number of squares!

Using conditional independence

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

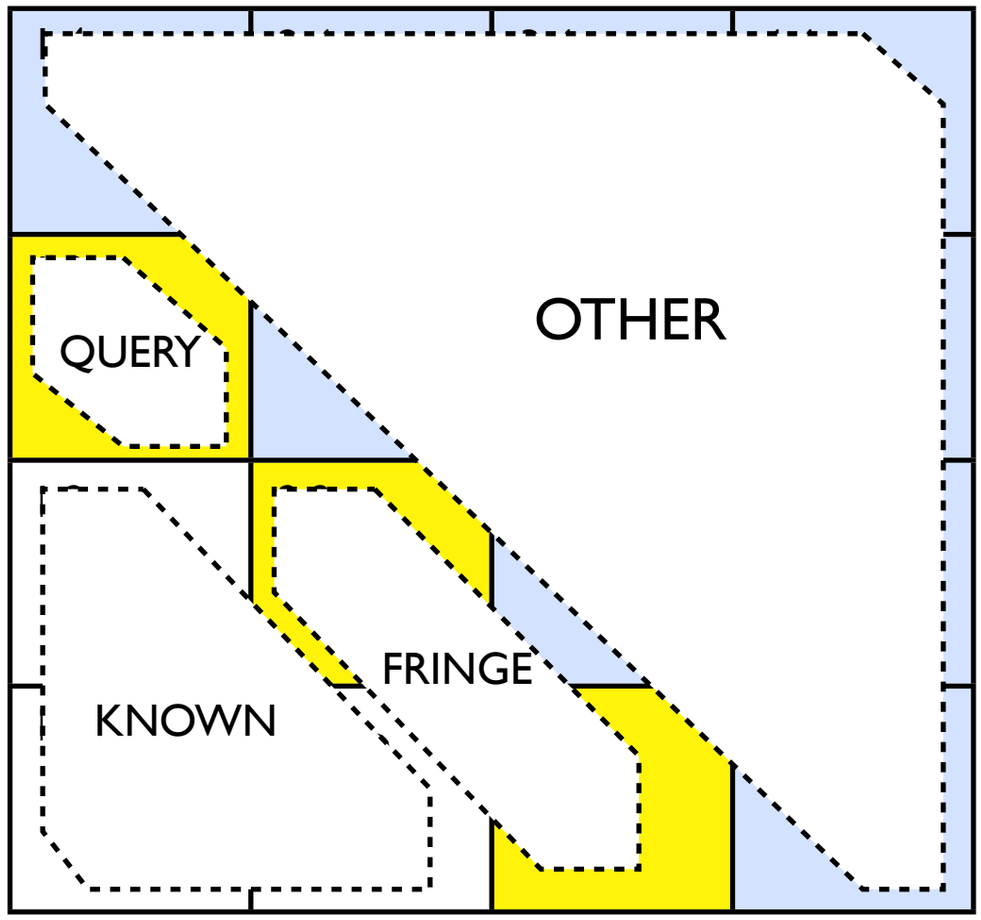

Define *Unknown = Fringe ∪ Other*

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

Using conditional independence (2)

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

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

$$= \alpha \sum_{fringe} \sum_{other} \mathbf{P}(b \mid known, P_{1,3}, fringe, other)\, \mathbf{P}(P_{1,3}, known, fringe, other)$$

$$= \alpha \sum_{fringe} \sum_{other} \mathbf{P}(b \mid known, P_{1,3}, fringe)\, \mathbf{P}(P_{1,3}, known, fringe, other)$$

$$= \alpha \sum_{fringe} \mathbf{P}(b \mid known, P_{1,3}, fringe) \sum_{other} \mathbf{P}(P_{1,3}, known, fringe, other)$$

$$= \alpha \sum_{fringe} \mathbf{P}(b \mid known, P_{1,3}, fringe) \sum_{other} \mathbf{P}(P_{1,3})\, P(known)\, P(fringe)\, P(other)$$

$$= \alpha\, P(known)\, \mathbf{P}(P_{1,3}) \sum_{fringe} \mathbf{P}(b \mid known, P_{1,3}, fringe)\, P(fringe) \sum_{other} P(other)$$

$$= \alpha' \mathbf{P}(P_{1,3}) \sum_{fringe} \mathbf{P}(b \mid known, P_{1,3}, fringe)\, P(fringe)$$

Wumpus World

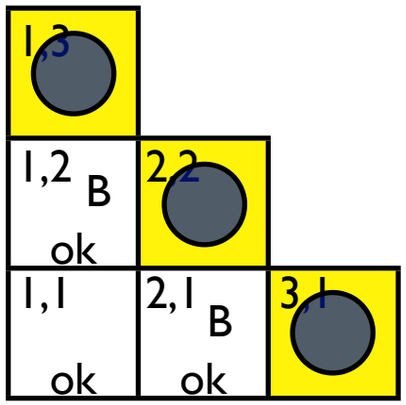

0.2 * 0.2 = 0.04

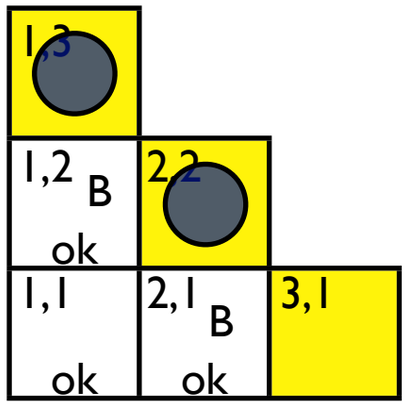

0.2 * 0.8 = 0.16

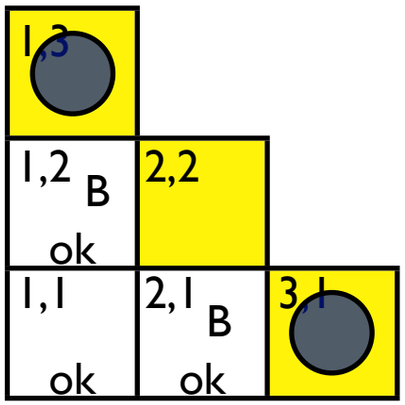

0.8 * 0.2 = 016

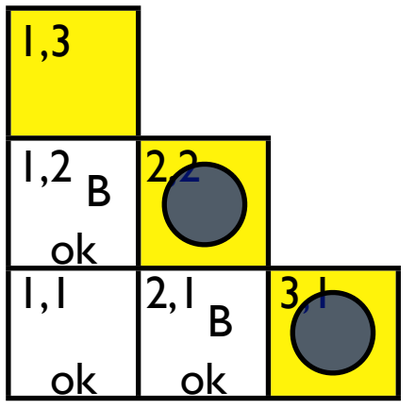

0.2 * 0.2 = 0.04

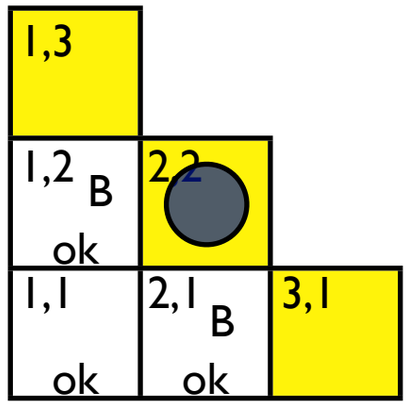

0.2 * 0.8 = 0.16

$$\boldsymbol{P}(P_{1,3} \mid known, b) = \alpha' \langle 0.2\,(0.04 + 0.16 + 0.16), 0.8\,(0.04 + 0.16) \rangle$$

$$\approx \langle 0.31, 0.69 \rangle$$

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

Summary

Probability is a way to formalise and represent uncertain knowledge

The *joint probability distribution* specifies probability over 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

Bayes' rule can be applied to compute posterior probabilities so that *diagnostic* probabilities can be assessed from *causal* ones

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Bayesian networks

A simple, graphical notation for conditional independence assertions and hence for compact specification of full joint distributions

Syntax:

a set of nodes, one per random variable

a directed, acyclic graph (link ≈ "directly influences")

a conditional distribution for each node given its parents:

$\boldsymbol{P}(X_i \mid Parents(X_i))$

In the simplest case, conditional distribution represented as a

conditional probability table (CPT)

giving the distribution over X_i for each combination of parent values

Example

Topology of network encodes conditional independence assertions:

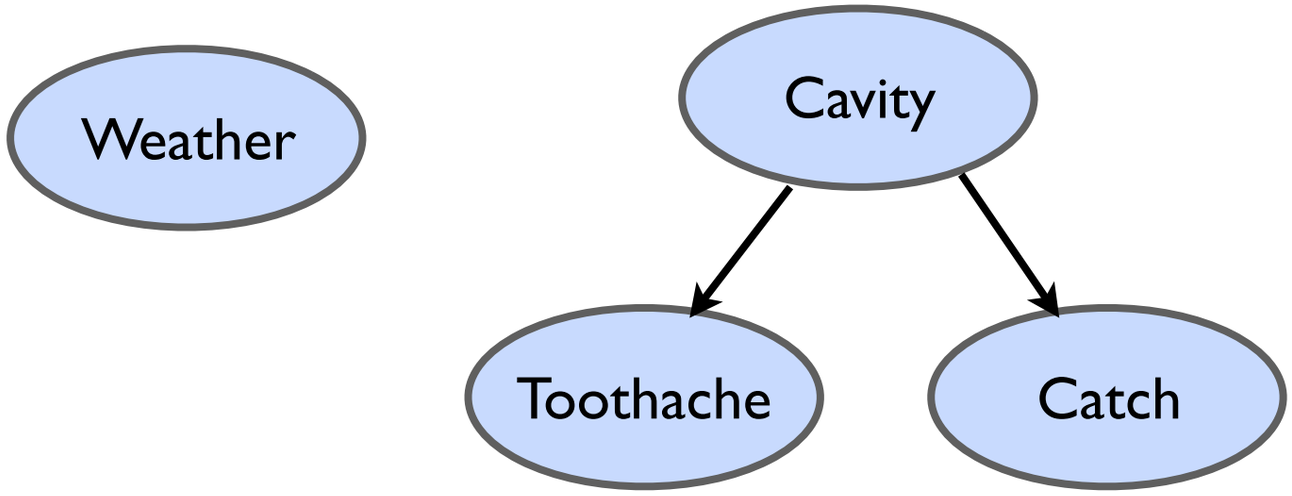

Weather is independent of the other variables

Toothache and *Catch* are conditionally independent given *Cavity*

Example 2

I am at work, my neighbour John calls to say my alarm is ringing, but neighbour Mary does not call.

Sometimes the alarm is set off by minor earthquakes.

Is there a burglar?

Variables: *Burglar*, *Earthquake*, *Alarm*, *JohnCalls*, *MaryCalls*

Network topology reflects "causal" knowledge:

A burglar can set the alarm off

An earthquake can set the alarm off

The alarm can cause John to call

The alarm can cause Mary to call

Example 2 (2)

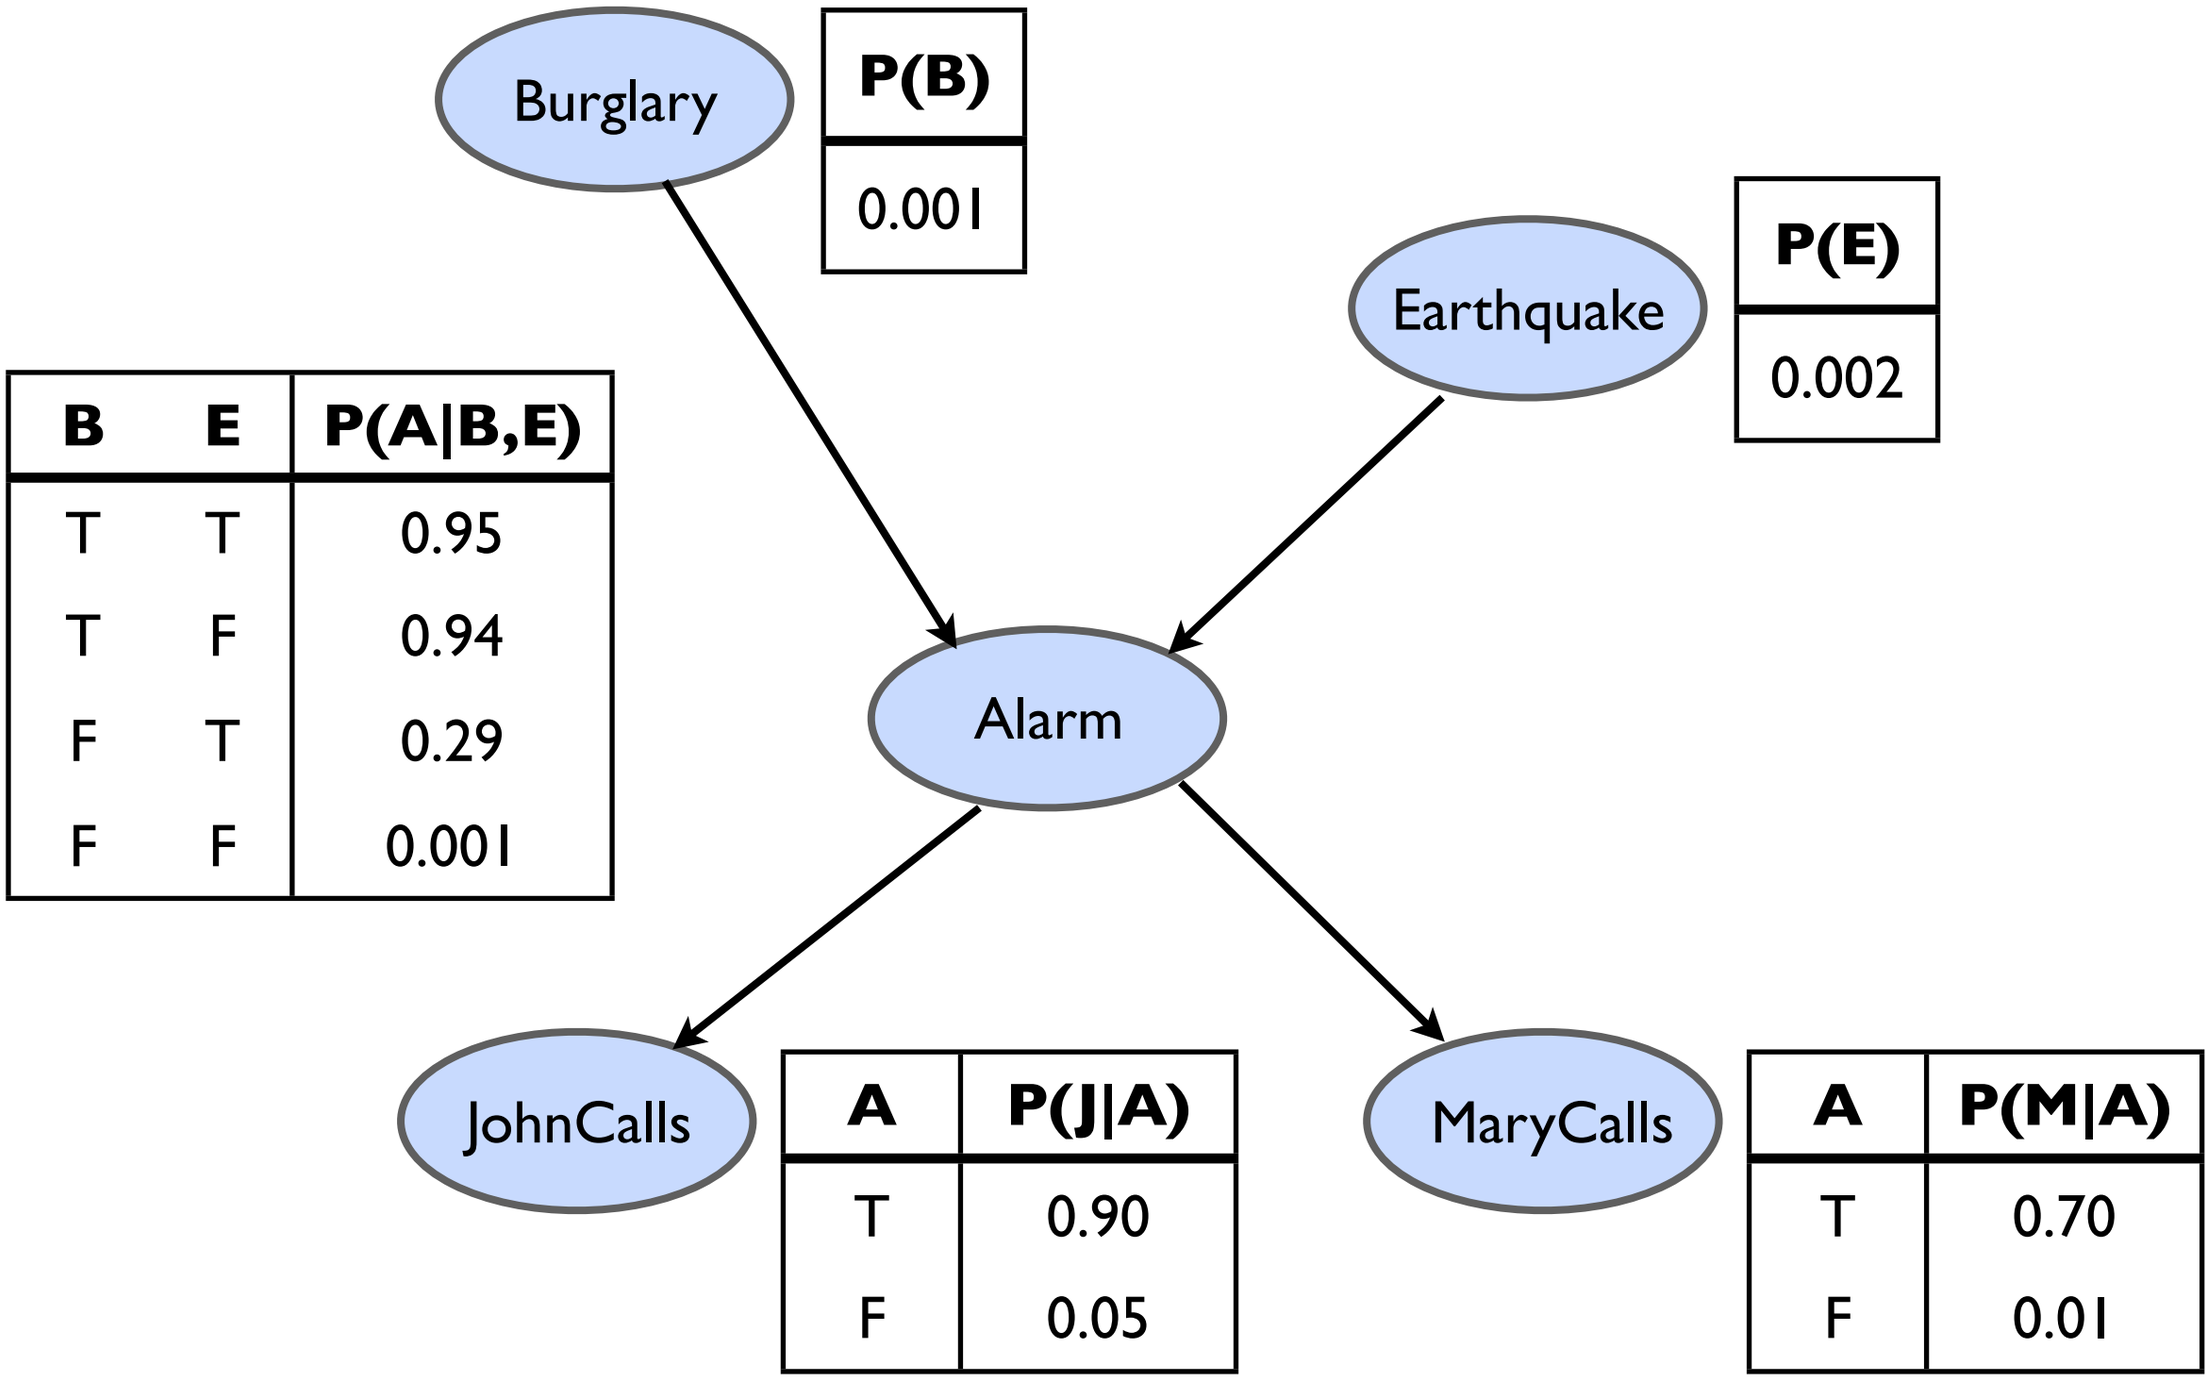

B	E	P(A|B,E)
T	T	0.95
T	F	0.94
F	T	0.29
F	F	0.001

A	P(J|A)
T	0.90
F	0.05

A	P(M|A)
T	0.70
F	0.01

Example 2

A CPT for Boolean X_i with k Boolean parents has 2^k rows for the combinations of parent values

Each row requires one number p for X_i = *true* (the number for X_i = *false* is just $1-p$)

If each variable has no more than k parents, the complete network requires $O(n\ 2^k)$ numbers

I.e., grows linearly with n, vs. $O(2^n)$ for the full joint distribution

For burglary net, $1 + 1 + 4 + 2 + 2 = 10$ numbers (vs. $2^5 - 1 = 31$)

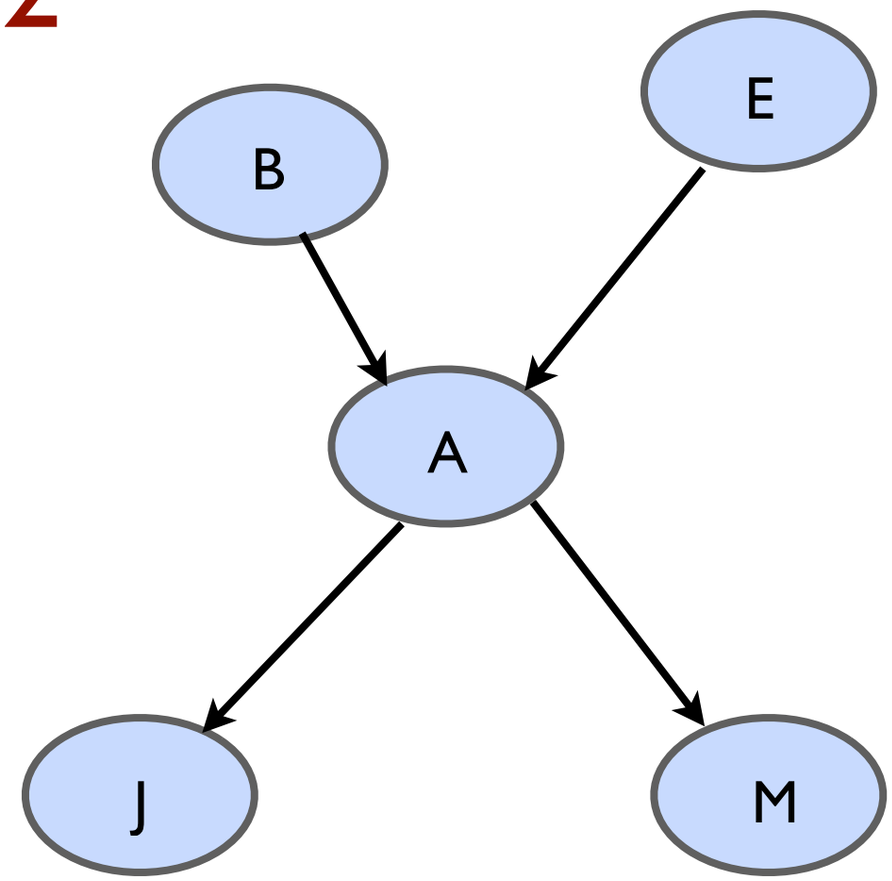

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Global semantics

Global semantics defines the full joint distribution as the product of the local conditional distributions:

$$P(x_1, \ldots, x_n) = \prod_{i=1}^{n} P(x_i \mid parents(X_i))$$

E.g., $P(j \wedge m \wedge a \wedge \neg b \wedge \neg e)$

=

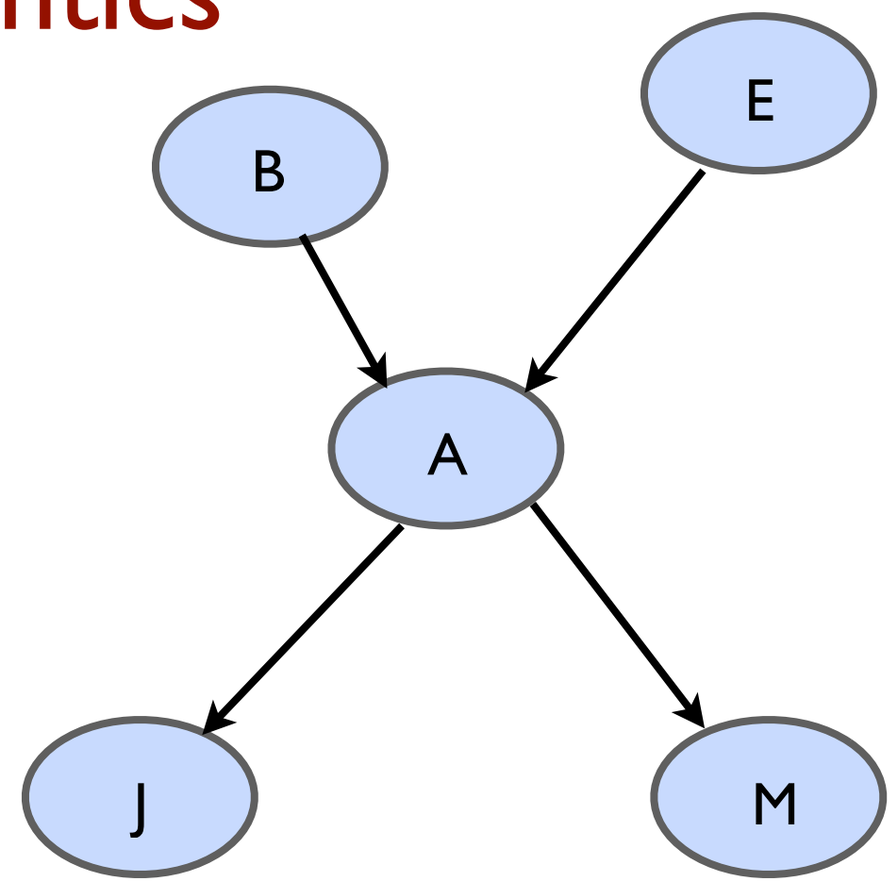

Global semantics

Global semantics defines the full joint distribution as the product of the local conditional distributions:

$$P(x_1, \ldots, x_n) = \prod_{i=1}^{n} P(x_i \mid parents(X_i))$$

E.g., $P(j \wedge m \wedge a \wedge \neg b \wedge \neg e)$

$$= P(j \mid a)\, P(m \mid a)\, P(a \mid \neg b, \neg e)\, P(\neg b)\, P(\neg e)$$

$$= 0.9 * 0.7 * 0.001 * 0.999 * 0.998$$

$$\approx 0.000628$$

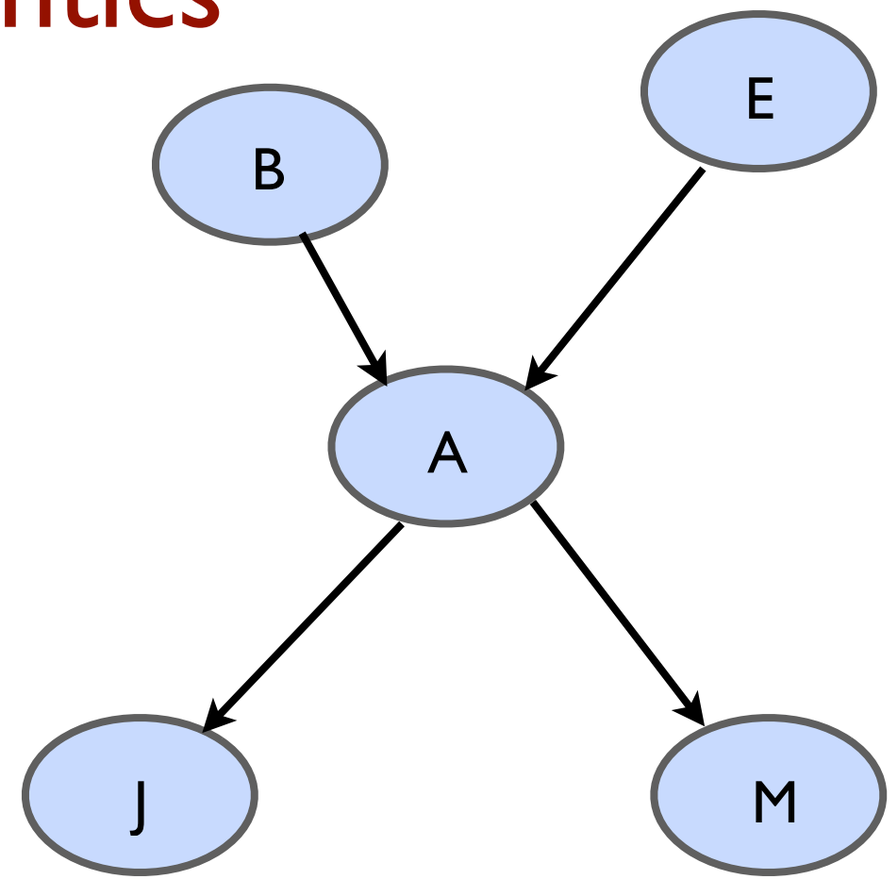

Constructing Bayesian networks

We need a method such that a series of locally testable assertions of conditional independence guarantees the required global semantics.

1. Choose an ordering of variables $X_1,..., X_n$

2. For $i = 1$ to n

    add X_i to the network

    select parents from $X_1,..., X_{i-1}$ such that

    $$\mathbf{P}(X_i \mid Parents(X_i)) = \mathbf{P}(X_i \mid X_1,..., X_{i-1})$$

This choice of parents guarantees the global semantics:

$$\mathbf{P}(X_1,..., X_n) = \prod_{i=1}^{n} \mathbf{P}(X_i \mid X_1,..., X_{i-1}) \quad \text{(chain rule)}$$

$$= \prod_{i=1}^{n} \mathbf{P}(X_i \mid Parents(X_i)) \quad \text{(by construction)}$$

Construction example

Suppose we choose the ordering M, J, A, B, E

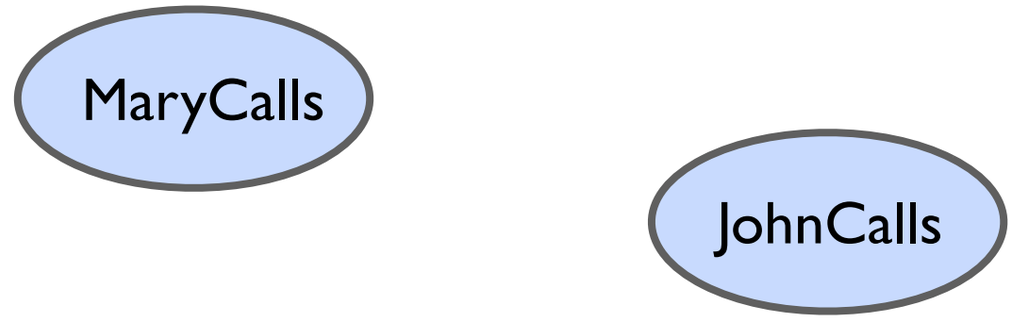

$P(J \mid M) = P(J)$?

Construction example

Suppose we choose the ordering M, J, A, B, E

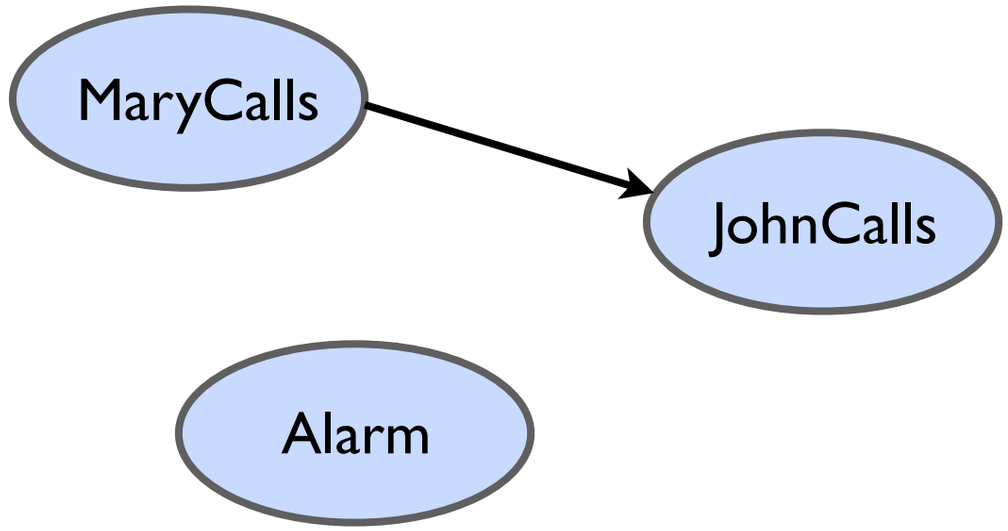

$P(J \mid M) = P(J)$? No

$P(A \mid J, M) = P(A \mid J)$? $P(A \mid J, M) = P(A)$?

Construction example

Suppose we choose the ordering M, J, A, B, E

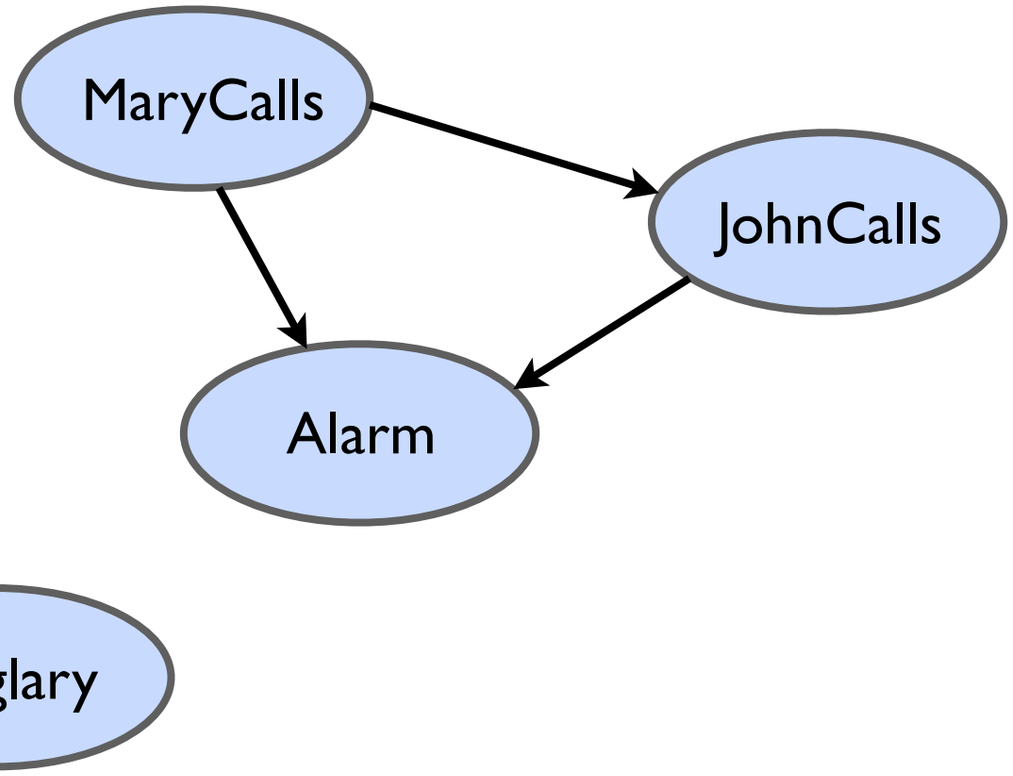

$P(J \mid M) = P(J)$? No

$P(A \mid J, M) = P(A \mid J)$? $P(A \mid J, M) = P(A)$? No

$P(B \mid A, J, M) = P(B \mid A)$?

$P(B \mid A, J, M) = P(B)$?

Construction example

Suppose we choose the ordering M, J, A, B, E

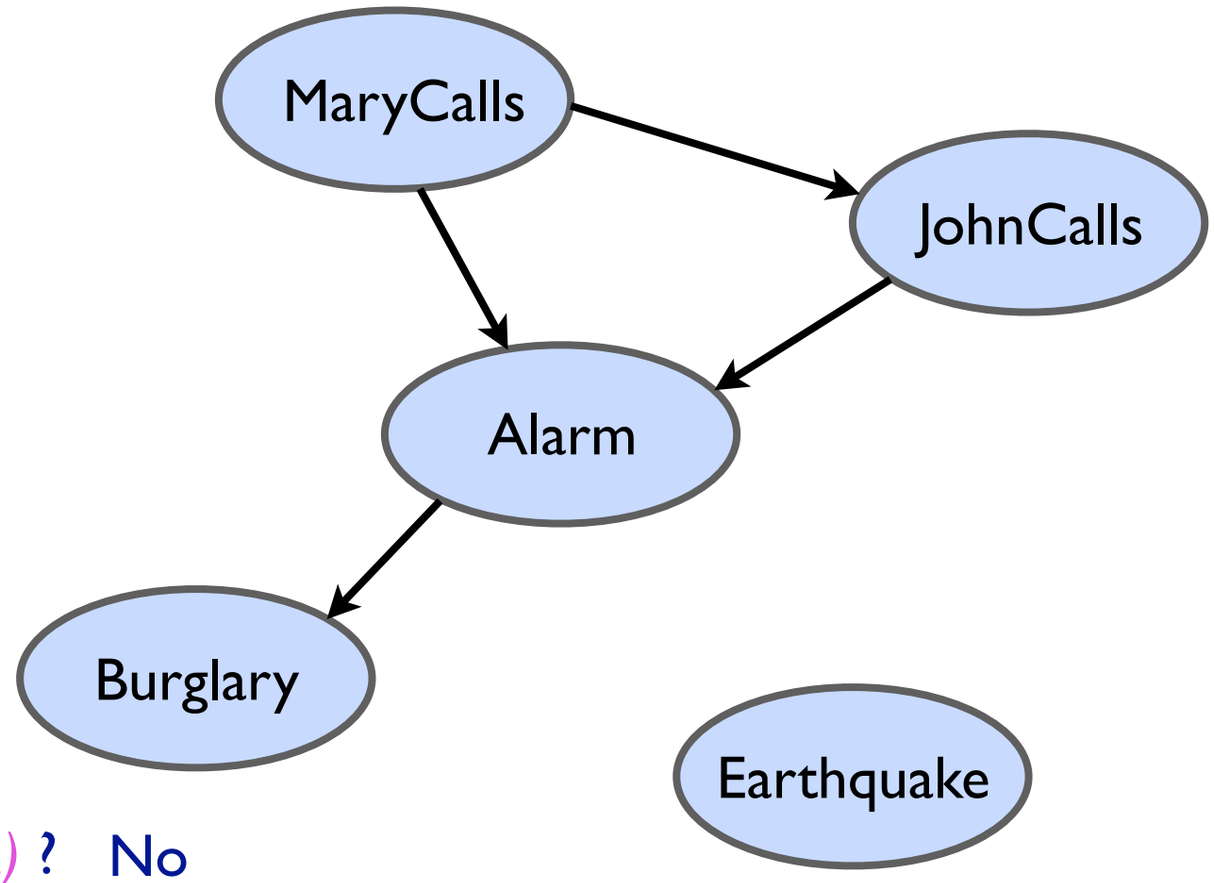

$P(J \mid M) = P(J)$? No

$P(A \mid J, M) = P(A \mid J)$? $P(A \mid J, M) = P(A)$? No

$P(B \mid A, J, M) = P(B \mid A)$? Yes

$P(B \mid A, J, M) = P(B)$? No

$P(E \mid B, A, J, M) = P(E \mid A)$?

$P(E \mid B, A, J, M) = P(E \mid A, B)$?

Construction example

Suppose we choose the ordering M, J, A, B, E

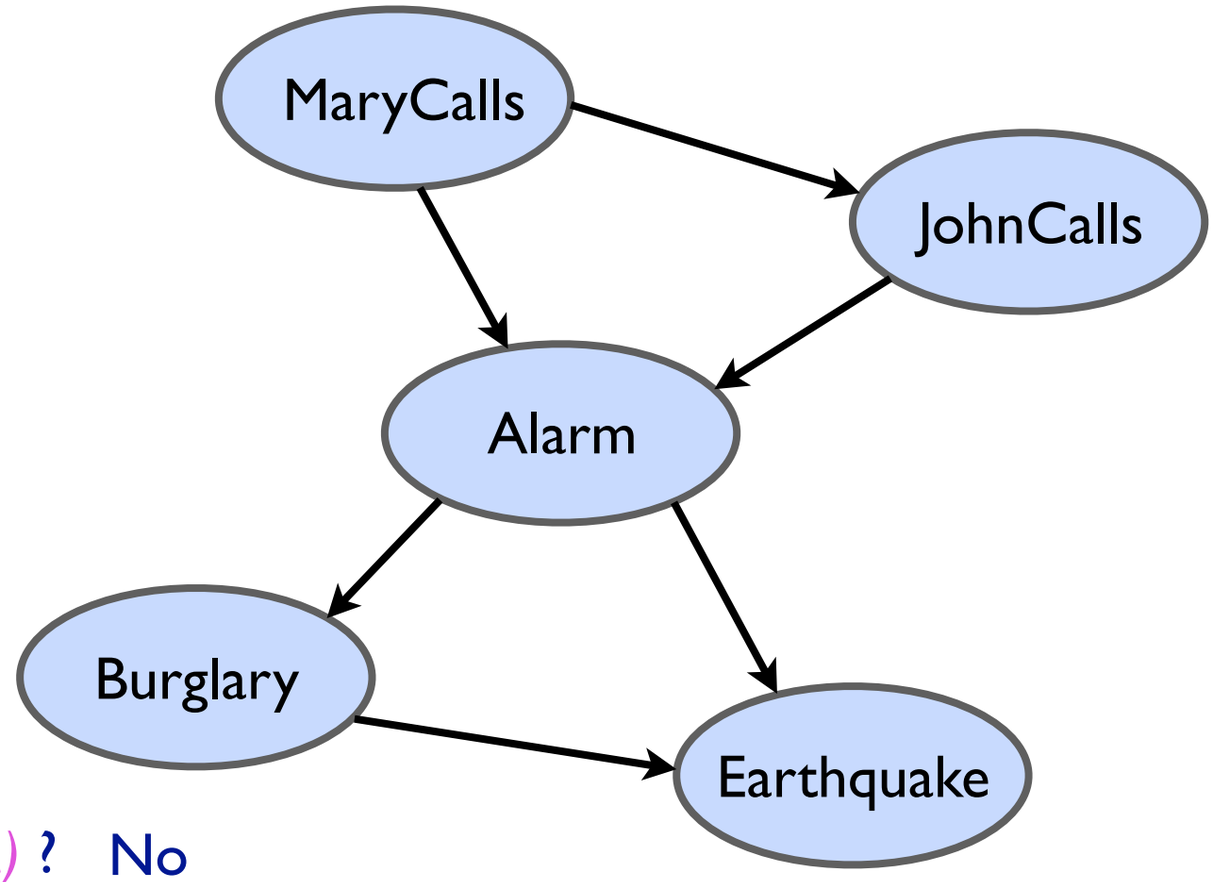

$P(J \mid M) = P(J)$? No

$P(A \mid J, M) = P(A \mid J)$? $P(A \mid J, M) = P(A)$? No

$P(B \mid A, J, M) = P(B \mid A)$? Yes

$P(B \mid A, J, M) = P(B)$? No

$P(E \mid B, A, J, M) = P(E \mid A)$? No

$P(E \mid B, A, J, M) = P(E \mid A, B)$? Yes

Construction example

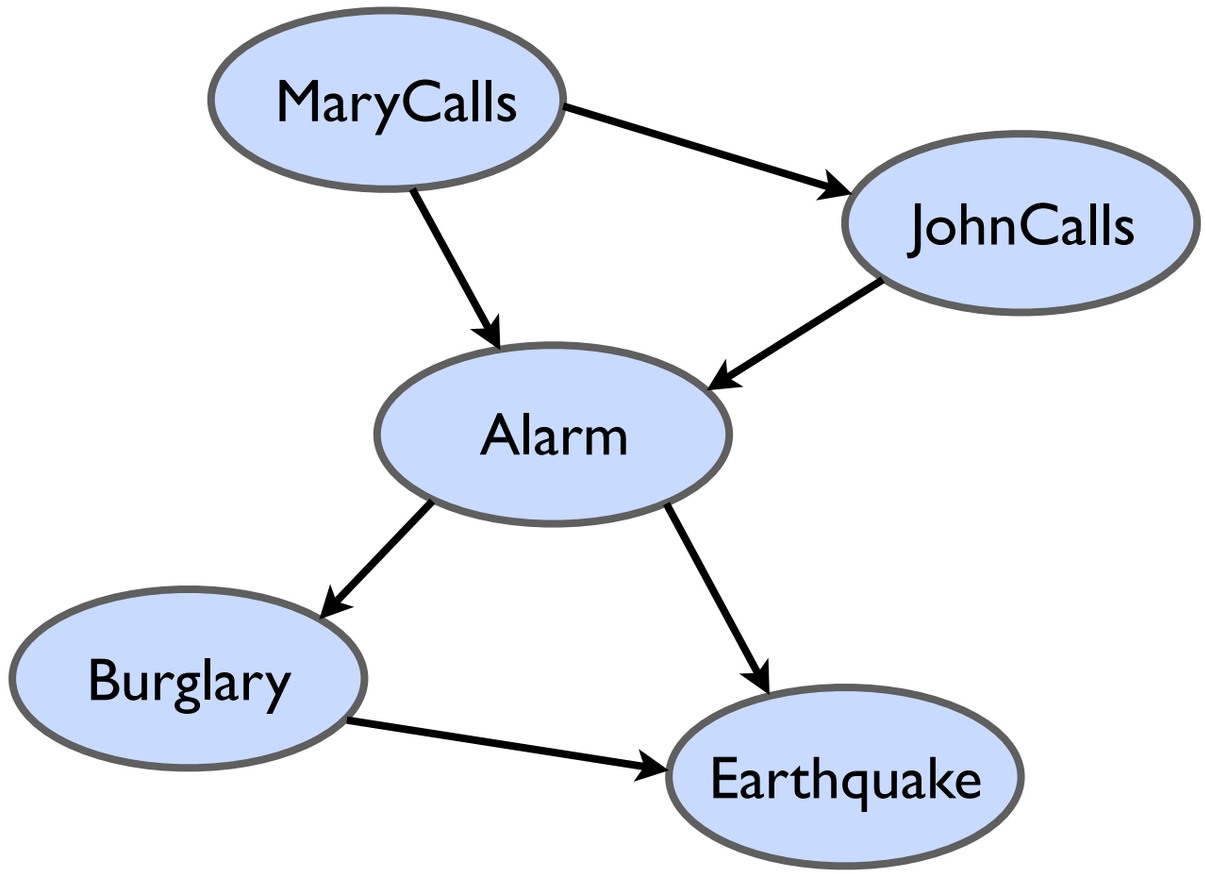

Deciding conditional independence is hard in noncausal directions

(Causal models and conditional independence seem hardwired for humans!)

Assessing conditional probabilities is hard in noncausal directions

Network is less compact: *1 + 2 + 4 +2 +4 = 13* numbers

Hence: Choose preferably an order corresponding to the cause → effect "chain"

Locally structured (sparse): Car diagnosis

Initial evidence: The *** car won't start!

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

Hidden variables (blue) ensure sparse structure / reduce parameters

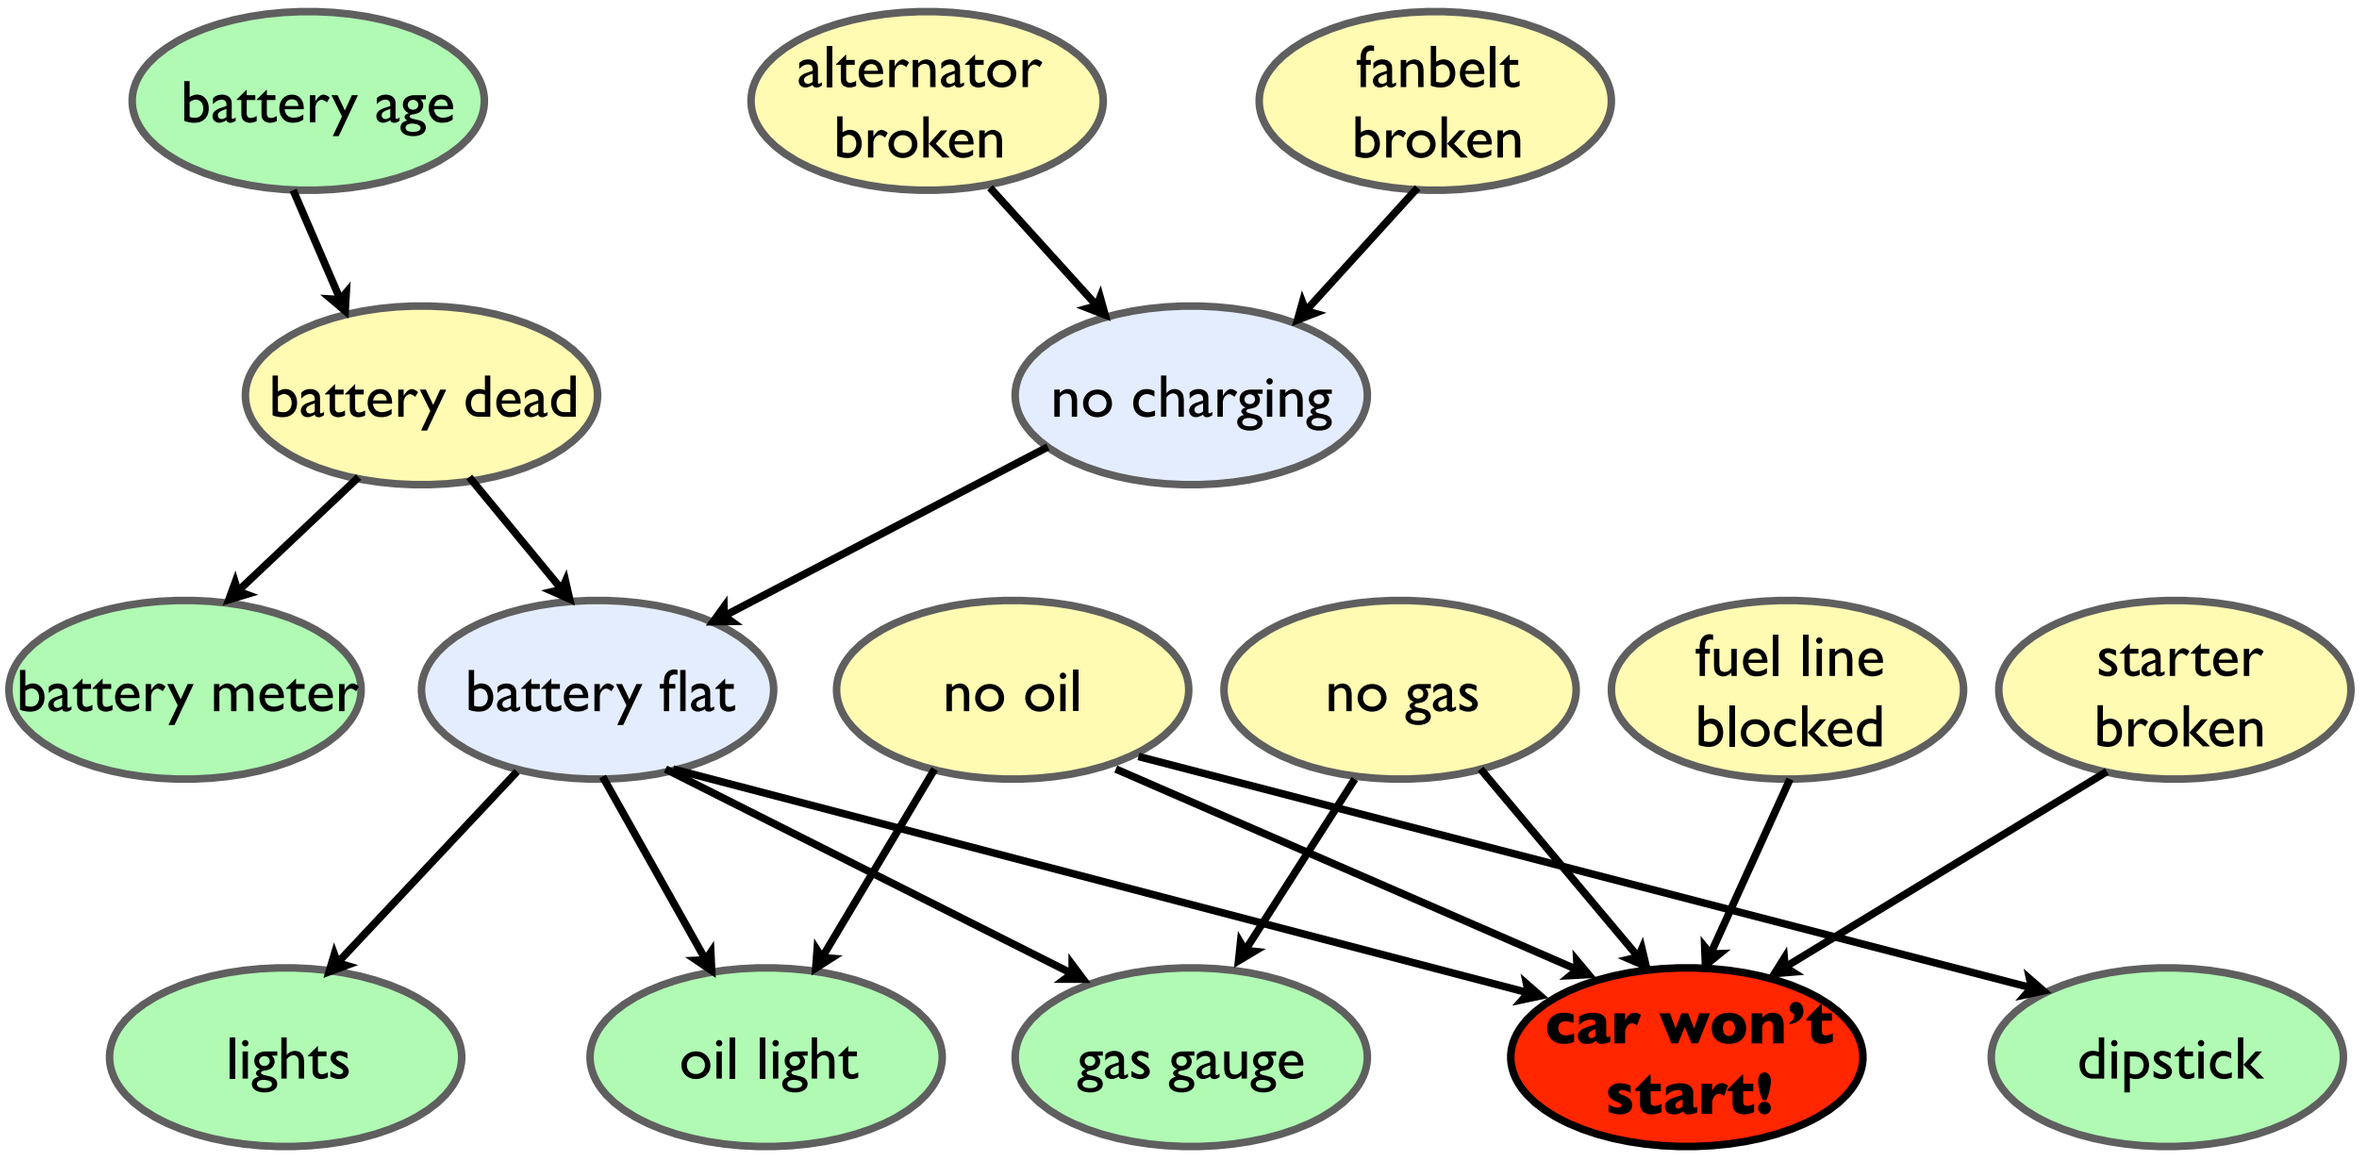

Local semantics

Local semantics: each node is conditionally independent of its non-descendants given its parents

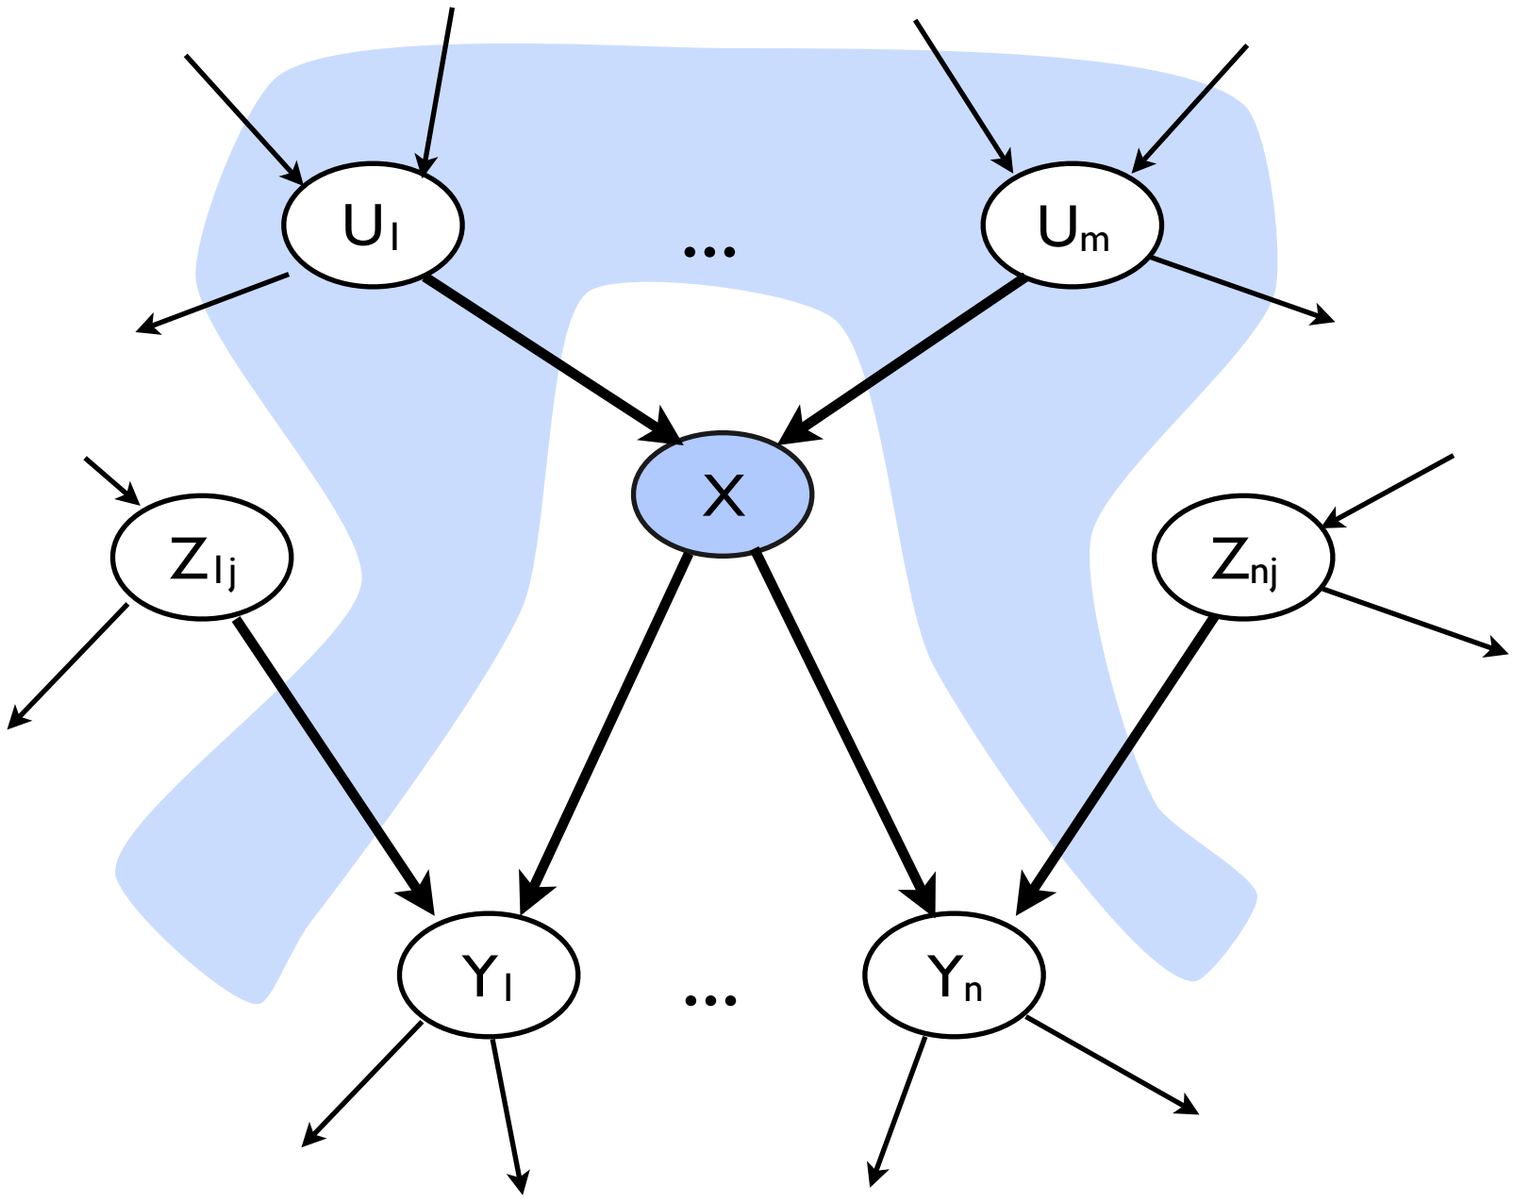

Markov blanket

Each node is conditionally independent of all others given its

Markov blanket: parents + children + children's parents

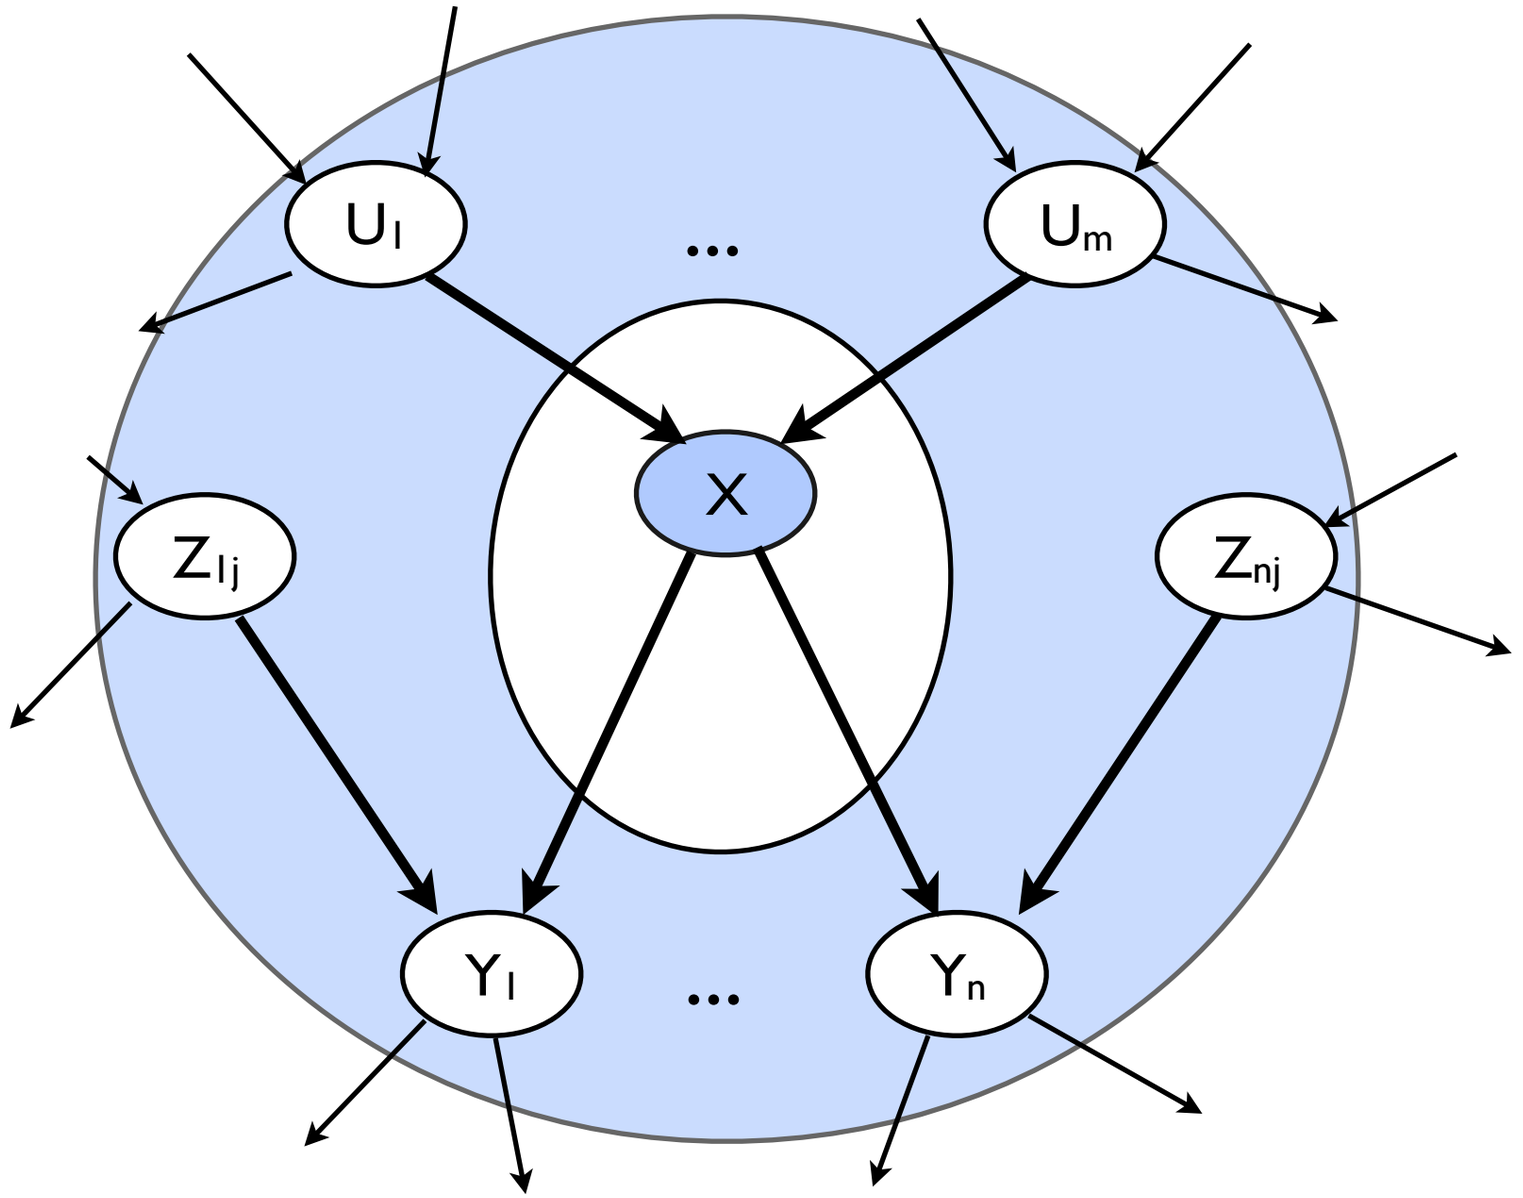

Outline

- Uncertainty (chapter 13)
  - Uncertainty
  - Probability
  - Syntax and Semantics
  - Inference
  - Independence and Bayes' Rule
- Bayesian Networks (chapter 14.1-3)
  - Syntax
  - Semantics
  - Efficient representation of conditional distributions (parameterised distributions)

Compact conditional distributions

CPT grows exponentially with numbers of parents (i.e., causes to the effect)

CPT becomes infinite with continuous-valued parent or child

Solution: *canonical* distributions that are defined compactly

Deterministic nodes are the simplest case:

$X = f(\mathit{Parents}(X))$ for some function f

E.g., Boolean functions

$$\mathit{NorthAmerican} \Leftrightarrow \mathit{Canadian} \vee \mathit{US} \vee \mathit{Mexican}$$

E.g., numerical relationships among continuous variables

$$\frac{\delta \mathit{Level}}{\delta t} = \mathit{inflow} + \mathit{precipitation} - \mathit{outflow} - \mathit{evaporation}$$

Compact conditional distributions (2)

Noisy-OR distributions model multiple noninteracting causes

1) Parents $U_1 \ldots U_k$ include all causes (add *leak node* for "miscellaneous" ones)

2) Independent failure probability q_i for each cause alone

$$\Rightarrow P(X \mid U_1, \ldots, U_j, \neg U_{j+1}, \ldots, \neg U_k) = 1 - \prod_{i=1}^{j} q_i$$

Cold	*Flu*	*Malaria*	*P(Fever)*	*P(¬Fever)*
F	F	F	***0.0***	*1.0*
F	F	T	*0.9*	***0.1***
F	T	F	*0.8*	***0.2***
F	T	T	*0.98*	*0.02 = 0.2 * 0.1*
T	F	F	*0.4*	***0.6***
T	F	T	*0.94*	*0.06 = 0.6 * 0.1*
T	T	F	*0.88*	*0.12 = 0.6 * 0.2*
T	T	T	*0.988*	*0.012 = 0.6 * 0.2 * 0.1*

Number of parameters linear in number of parents

Hybrid (discrete + continuous) networks

Discrete (*Subsidy?* and *Buys?*); continuous (*Harvest* and *Cost*)

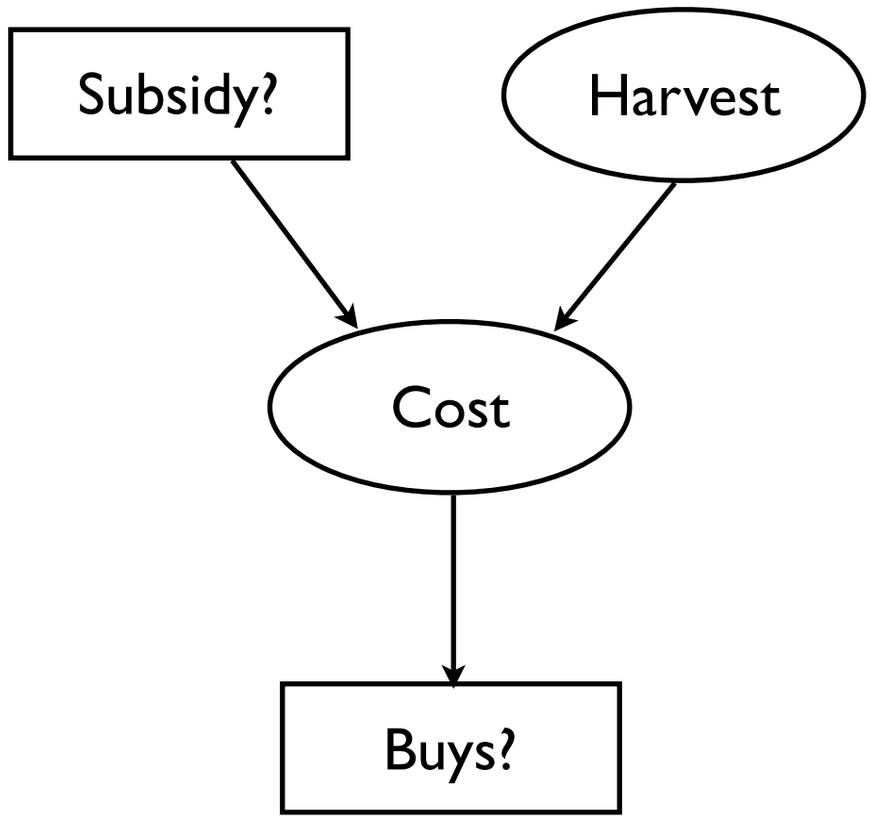

Option 1: discretisation - possibly large errors, large CPTs

Option 2: finitely parameterised canonical families

1) Continuous variable, discrete + continuous parents (e.g., *Cost*)

2) Discrete variable, continuous parents (e.g., *Buys?*)

Continuous child variables

Need one *conditional density* function for child variable given continuous parents, for each possible assignment to discrete parents

Most common is the *linear Gaussian* model (following *G. normal distribution*), e.g.:

$$P(\mathit{Cost} = c \mid \mathit{Harvest} = h, \mathit{Subsidy?} = \mathit{true})$$

$$= N(a_t h + b_t, \sigma_t)(c)$$

Mean *Cost* varies linearly with *Harvest*, variance is fixed

Linear variation is unreasonable over the full range

but works OK if the likely range of *Harvest* is narrow

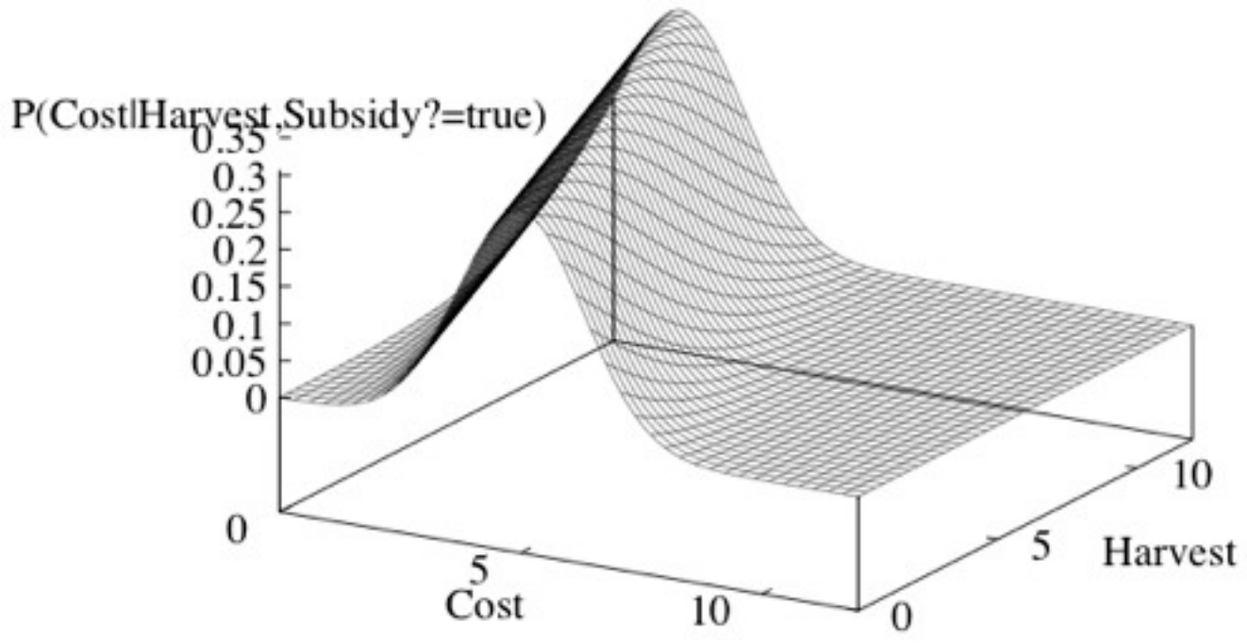

Discrete variable with continuous parents

Probability of *Buys?* given *Cost* should be a soft threshold

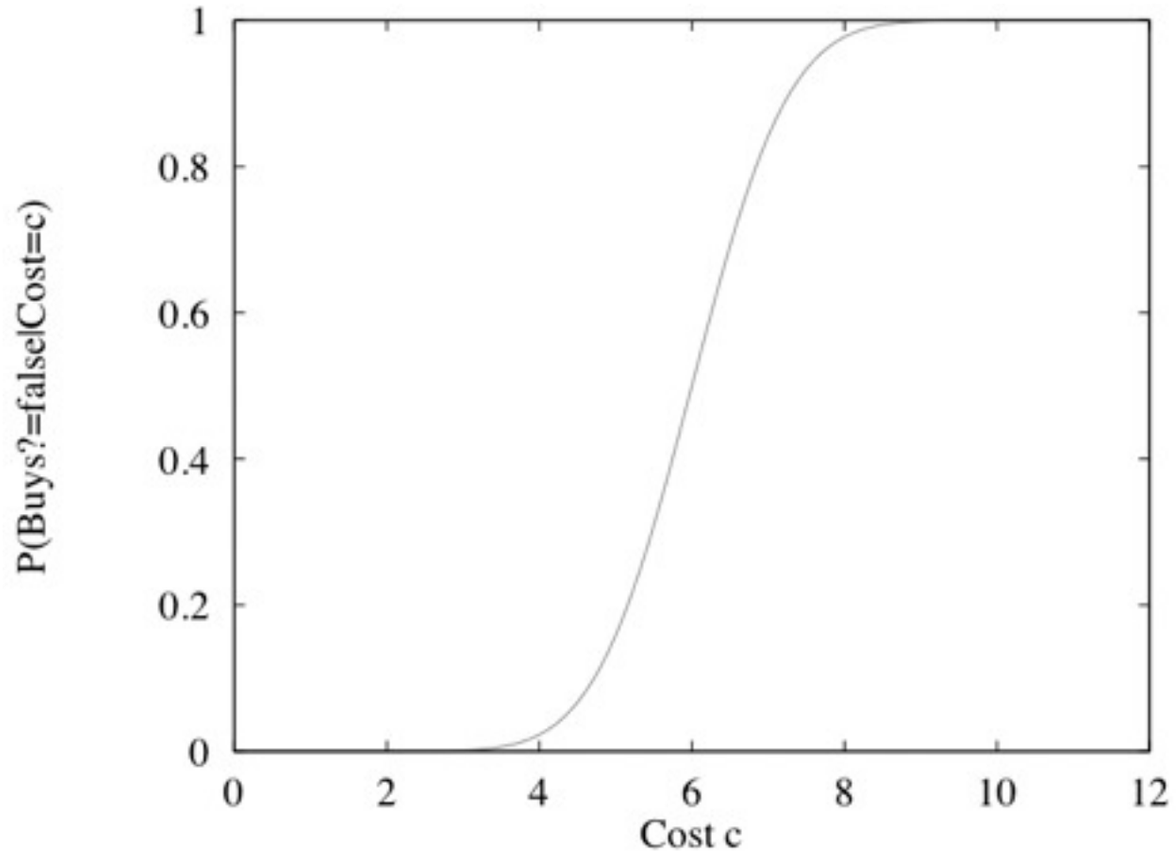

E.g., Probit (probability unit), the integral of the standard normal distribution, or "logit", the logistic function, also representing a type of sigmoid.

Summary

Bayesian networks provide a natural representation for (causally induced) conditional independence

Topology + CPTs = compact representation of joint distribution

Generally easy for (non)experts to construct

Canonical distributions (e.g., noisy-OR) = compact representation of CPTs

Continuous variables ⇒ parameterised distributions (e.g., linear Gaussians)