Reasoning with Graphical Models

Slides Set 2:

Rina Dechter

Reading:
Darwiche chapter 4
Pearl: chapter 3

Outline

- Bayesian Networks, DAGS, Markov(G)
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Outline

- Bayesian Networks, DAGS, Markov(G)
  - From a distribution to a BN
  - From BN to distributions, DAGs, Markov(G)
  - Parameterization
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Outline

- Bayesian Networks, DAGS, Markov(G)
  - From a distribution to a BN
  - From BN to distributions, DAGs, Markov(G)
  - Parameterization
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Bayesian Networks (BNs) in 2 ways:

From a distribution to a BN:

- A Bayesian network is factorize probability distribution along an ordering.
- The DAG emerging is a Bayesian network of the distribution
- The factorization is guided by a set of Markov assumption that transform the chain product formula into a Bayesian network.

From a BN to a distribution:

- Generate a DAG with its Markov assumptions.
- Parameterize the DAG yielding a Bayesian network which corresponds to a single probability distribution obtained by product.

- The BN distribution obeys additional independence assumption read from the DAG and can be proved using the Graphoid axioms.

Difficulty: Complexity in model construction and inference

- In Alarm example:
  - 31 numbers needed,
  - Quite unnatural to assess: e.g.

$$P(B=y, E=y, A=y, J=y, M=y)$$

  - Computing $P(B=y|M=y)$ takes 29 additions.
- In general,
  - $P(X_1, X_2, \ldots, X_n)$ needs at least $2^n - 1$ numbers to specify the joint probability. Exponential model size.
  - Knowledge acquisition difficult (complex, unnatural),
  - Exponential storage and inference.

Chain Rule and Factorization

Overcome the problem of exponential size by exploiting conditional independence

- The chain rule of probabilities:

$$
\begin{aligned}
P(X_1, X_2) &= P(X_1)P(X_2|X_1) \\
P(X_1, X_2, X_3) &= P(X_1)P(X_2|X_1)P(X_3|X_1, X_2) \\
\cdots & \\
P(X_1, X_2, \ldots, X_n) &= P(X_1)P(X_2|X_1)\ldots P(X_n|X_1, \ldots, X_{n-1}) \\
&= \prod_{i=1}^{n} P(X_i|X_1, \ldots, X_{i-1}).
\end{aligned}
$$

- No gains yet. The number of parameters required by the factors is: $2^{n-1} + 2^{n-1} + \ldots + 1 = 2^n - 1$.

Conditional Independence

- About $P(X_i|X_1, \ldots, X_{i-1})$:
  - Domain knowledge usually allows one to identify a subset $pa(X_i) \subseteq \{X_1, \ldots, X_{i-1}\}$ such that
    - Given $pa(X_i)$, X_i is independent of all variables in $\{X_1, \ldots, X_{i-1}\} \setminus pa(X_i)$, i.e.

$$P(X_i|X_1, \ldots, X_{i-1}) = P(X_i|pa(X_i))$$

- Then

$$P(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} P(X_i|pa(X_i))$$

- Joint distribution factorized.
- The number of parameters might have been substantially reduced.

Example continued

P(B,E,A,J,M)=?

P(B)P(E|B)P(A|B,E)P(J|B,E,A)P(M|B,E,A,J) =

P(B)P(E|B)P(A|B,E)P(J|A)P(M|A) =

pa(B) = {}, pa(E)={B},P(A)= {B,E}, pa(J) = {A}, pa(M) = {A}

Example continued

P(B,E,A,J,M)=?

P(B)P(E|B)P(A|B,E)P(J|B,E,A)P(M|B,E,A,J) =

P(B)P(E|B)P(A|B,E)P(J|A)P(M|A) =

pa(B) = {}, pa(E)={B},P(A)= {B,E}, pa(J) = {A}, pa(M) = {A}

- Conditional probabilities tables (CPT)

B	P(B)
Y	.01
N	.99

E	P(E)
Y	.02
N	.98

A	B	E	P(A|B, E)
Y	Y	Y	.95
N	Y	Y	.05
Y	Y	N	.94
N	Y	N	.06
Y	N	Y	.29
N	N	Y	.71
Y	N	N	.001
N	N	N	.999

M	A	P(M|A)
Y	Y	.9
N	Y	.1
Y	N	.05
N	N	.95

J	A	P(J|A)
Y	Y	.7
N	Y	.3
Y	N	.01
N	N	.99

Example continued

- Model size reduced from 31 to 1+1+4+2+2=10
- Model construction easier
  - Fewer parameters to assess.
  - Parameters more natural to assess:e.g.

$$P(B = Y), P(E = Y), P(A = Y|B = Y, E = Y),$$

$$P(J = Y|A = Y), P(M = Y|A = Y)$$

- Inference easier.Will see this later.

From Factorizations to Bayesian Networks

Graphically represent the conditional independency relationships:

- construct a directed graph by drawing an arc from X_j to X_i iff $X_j \in pa(X_i)$

$$pa(B) = \{\}, \; pa(E) = \{\}, \; pa(A) = \{B, E\}, \; pa(J) = \{A\}, \; pa(M) = \{A\}.$$

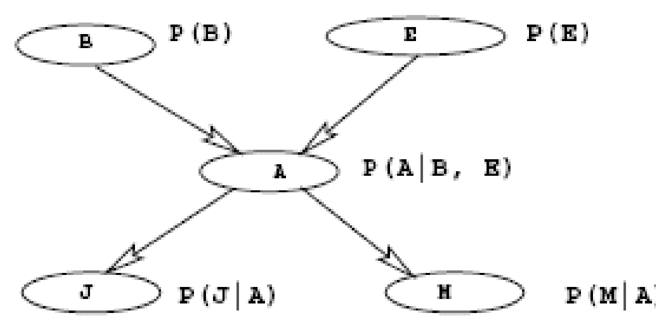

- Also attach the conditional probability (table) $P(X_i|pa(X_i))$ to node X_i.
- What results in is a Bayesian network. Also known as **belief network, probabilistic network.**

Outline

- Bayesian Networks, DAGS, Markov(G)
  - From a distribution to a BN
  - From BN to distributions, DAGs, Markov(G)
  - Parameterization
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Capturing Independence Graphically

The causal interpretation

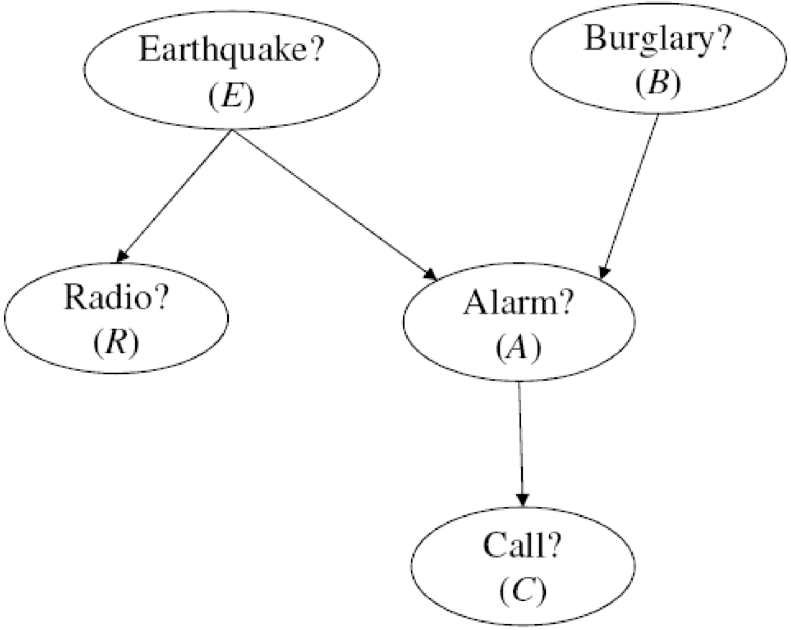

Assume that edges in this graph represent direct causal influences among these variables.

Example

The alarm triggering (A) is a direct cause of receiving a call from a neighbor (C).

Capturing Independence Graphically

But influences can be indirect as well.
For example…

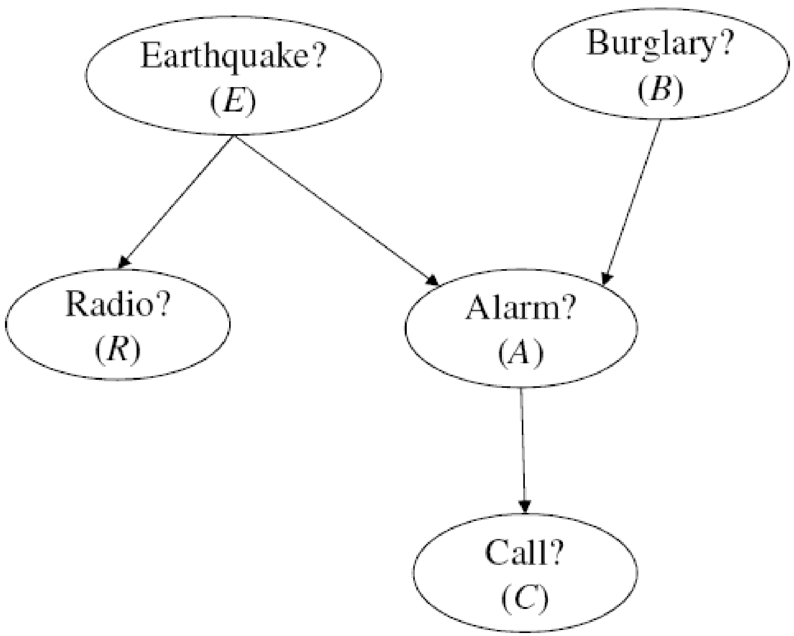

We expect our belief in C to be influenced by evidence on R.

Example

If we get a radio report that an earthquake took place in our neighborhood, our belief in the alarm triggering would probably increase, which would also increase our belief in receiving a call from our neighbor.

Capturing Independence Graphically

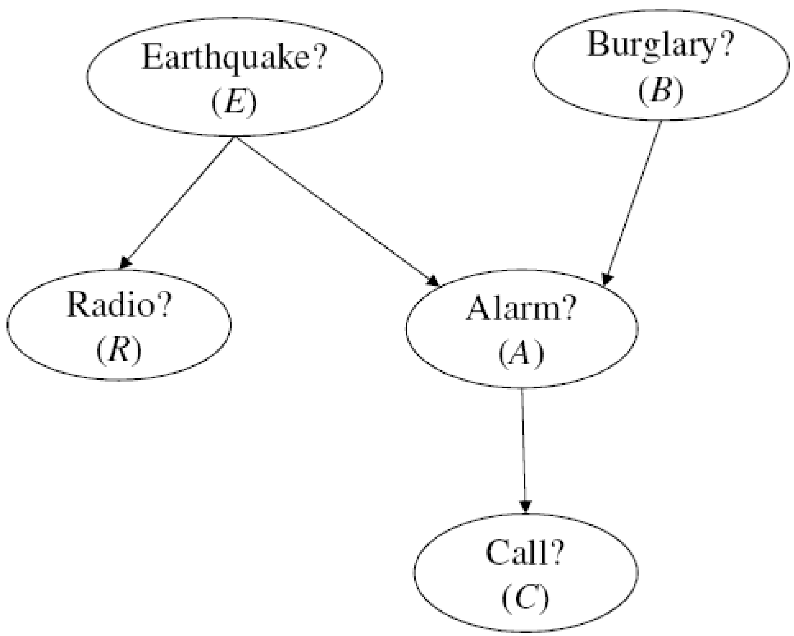

We would not change this belief, however, if we knew for sure that the alarm did not trigger. That is, we would find C independent of R given $\neg A$ in the context of this causal structure.

Capturing Independence Graphically

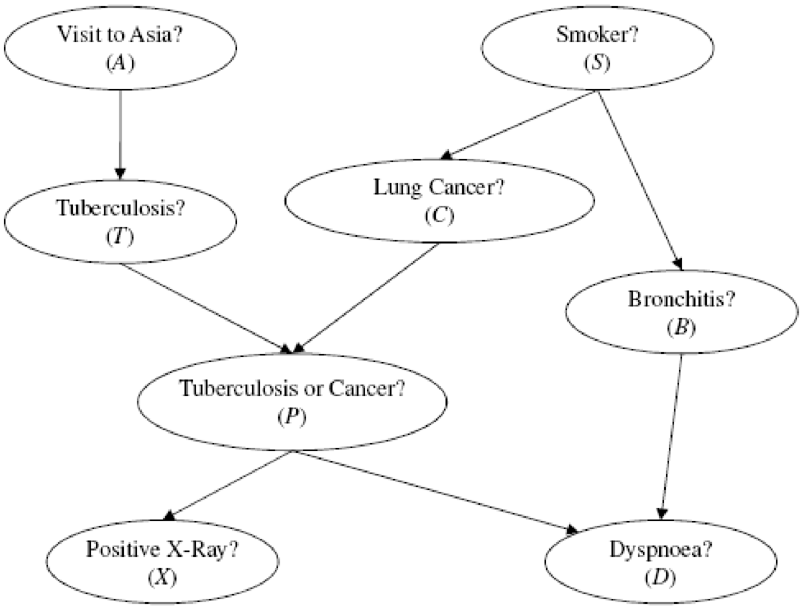

We would clearly find a visit to Asia relevant to our belief in the X-Ray test coming out positive, but we would find the visit irrelevant if we know for sure that the patient does not have Tuberculosis. That is, X is dependent on A, but is independent of A given $\neg T$.

Graphs Convey Independence Statements

- Directed graphs by graph’s d-separation
- Undirected graphs by graph separation
- Goal: capture probabilistic conditional independence by graphs.
- We focus on directed graphs first.

Capturing Independence Graphically

These examples of independence are all implied by a formal interpretation of each DAG as a set of conditional independence statements.

Given a variable V in a DAG G:

$\mathrm{Parents}(V)$ are the parents of V in DAG G, that is, the set of variables N with an edge from N to V.

$\mathrm{Descendants}(V)$ are the descendants of V in DAG G, that is, the set of variables N with a directed path from V to N (we also say that V is an ancestor of N in this case).

$\mathrm{Non\_Descendants}(V)$ are all variables in DAG G other than V, $\mathrm{Parents}(V)$ and $\mathrm{Descendants}(V)$. We will call these variables the non-descendants of V in DAG G.

Capturing Independence Graphically

We will formally interpret each DAG G as a compact representation of the following independence statements (Markovian assumptions):

$$I(V, \mathrm{Parents}(V), \mathrm{Non\_Descendants}(V)),$$

for all variables V in DAG G.

- If we view the DAG as a causal structure, then $\mathrm{Parents}(V)$ denotes the direct causes of V and $\mathrm{Descendants}(V)$ denotes the effects of V.
- Given the direct causes of a variable, our beliefs in that variable will no longer be influenced by any other variable except possibly by its effects.

Capturing Independence Graphically

What are the Markov assumptions here?

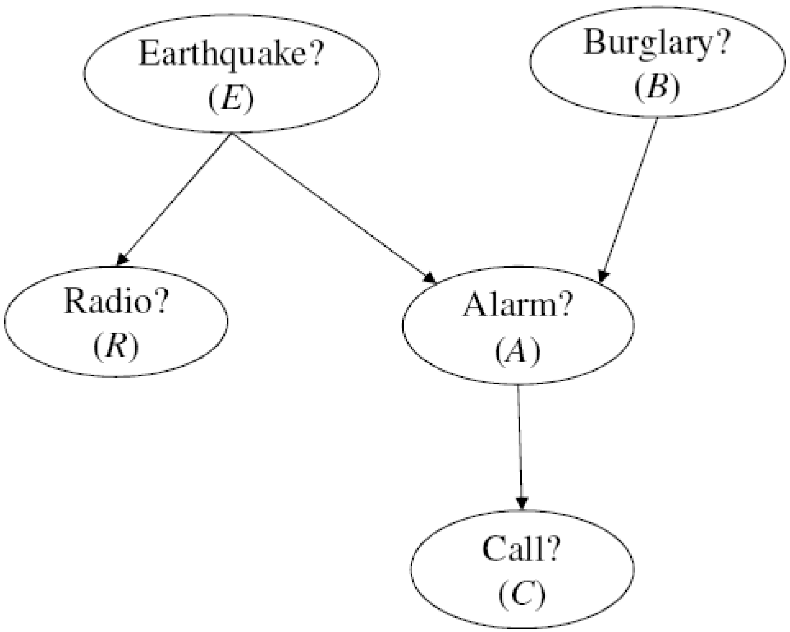

Note that variables B and E have no parents, hence, they are marginally independent of their non-descendants.

Capturing Independence Graphically

What are the Markov assumptions here?

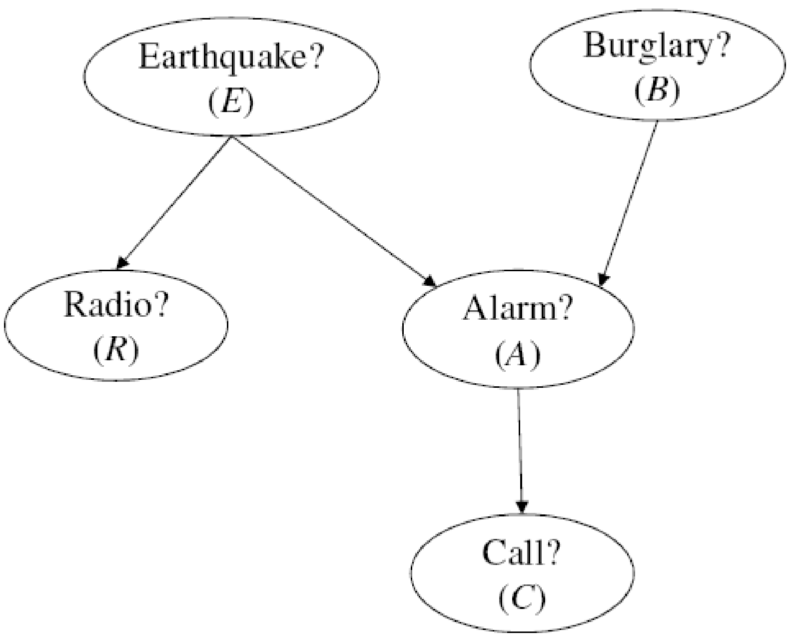

$I(C, A, \{B, E, R\})$
$I(R, E, \{A, B, C\})$
$I(A, \{B, E\}, R)$
$I(B, \emptyset, \{E, R\})$
$I(E, \emptyset, B)$

Note that variables B and E have no parents, hence, they are marginally independent of their non-descendants.

Capturing Independence Graphically

The formal interpretation of a DAG as a set of conditional independence statements makes no reference to the notion of causality, even though we have used causality to motivate this interpretation.

If one constructs the DAG based on causal perceptions, then one would tend to agree with the independencies declared by the DAG.

It is perfectly possible to have a DAG that does not match our causal perceptions, yet we agree with the independencies declared by the DAG.

Capturing Independence Graphically

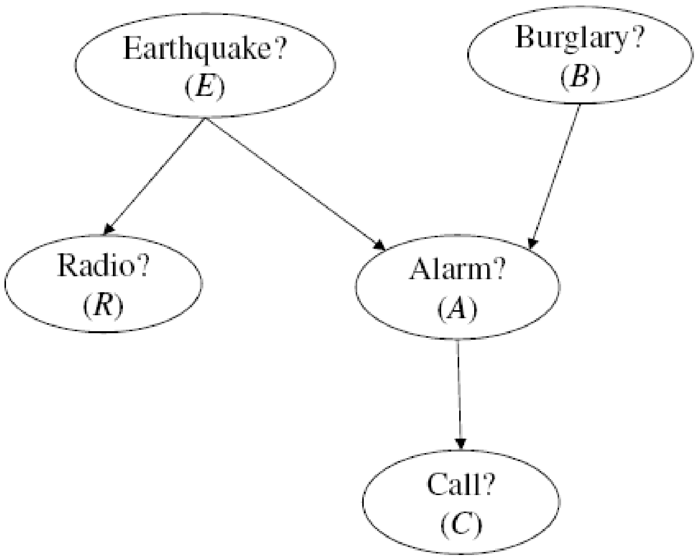

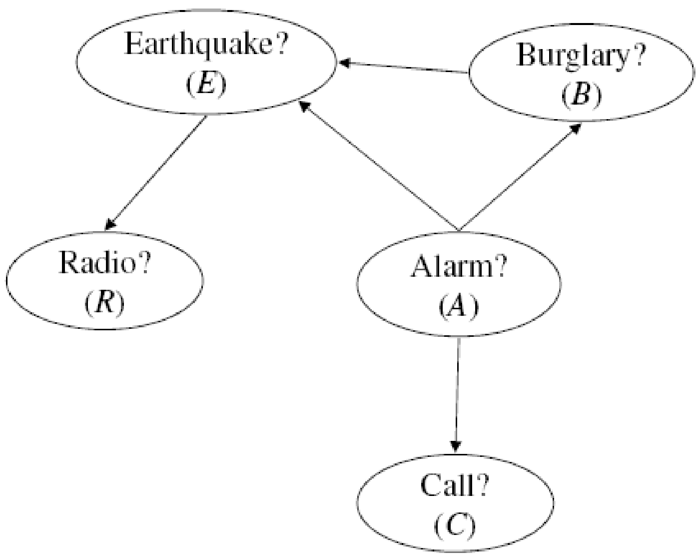

Every independence which is declared (or implied) by the second DAG is also declared (or implied) by the first one. Hence, if we accept the first DAG, then we must also accept the second.

Outline

- Bayesian Networks, DAGS, Markov(G)
  - From a distribution to a BN
  - From BN to distributions, DAGs, Markov(G)
  - Parameterization
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Parameterizing the Independence Structure

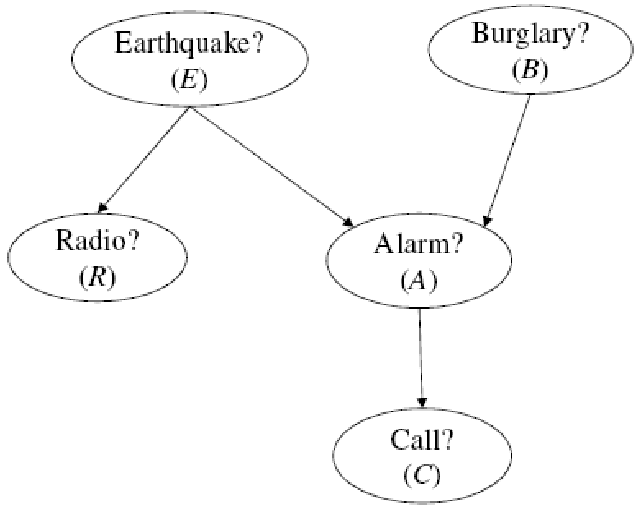

- The DAG G is a partial specification of our state of belief Pr.
- By constructing G, we are saying that the distribution Pr must satisfy the independence assumptions in $\mathrm{Markov}(G)$.
- This clearly constrains the possible choices for the distribution Pr, but does not uniquely define it.

We can augment the DAG G by a set of conditional probabilities that together with $\mathrm{Markov}(G)$ are guaranteed to define the distribution Pr uniquely.

Parameterizing the Independence Structure

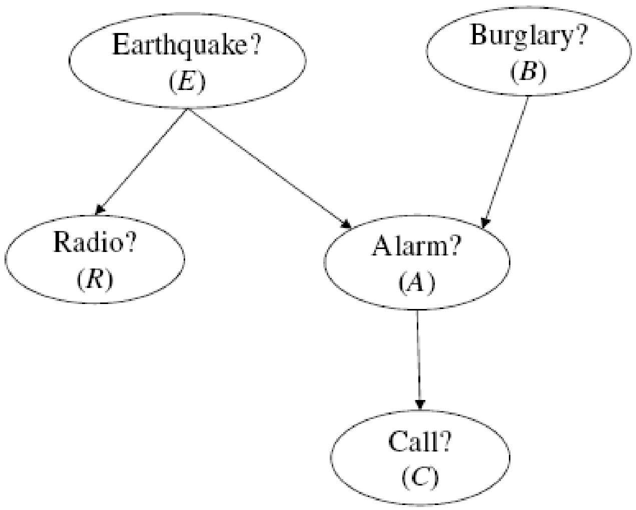

For every variable X in the DAG G, and its parents $\mathbf{U}$, we need to provide the probability $\Pr(x|\mathbf{u})$ for every value x of variable X and every instantiation $\mathbf{u}$ of parents $\mathbf{U}$.

Example

We need to provide the following conditional probabilities:

$$\Pr(c|a),\ \ \Pr(r|e),\ \ \Pr(a|b,e),\ \ \Pr(e),\ \ \Pr(b),$$

where a, b, c, e and r are values of variables A, B, C, E and R.

Parameterizing the Independence Structure

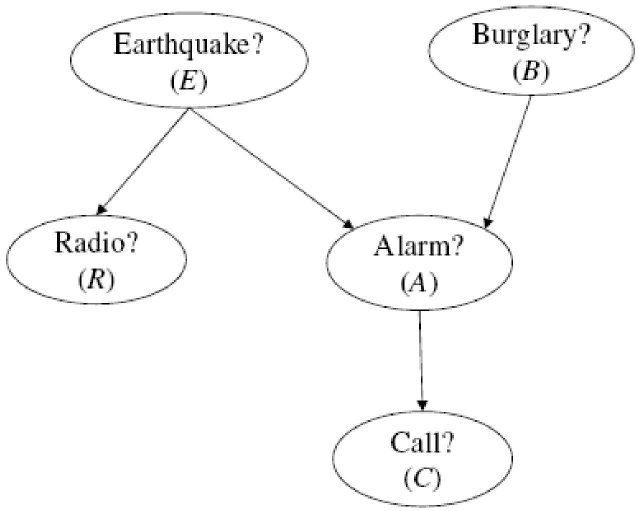

The conditional probabilities required for variable C:

A	C	$\Pr(c \mid a)$
true	true	.80
true	false	.20
false	true	.001
false	false	.999

The above table is known as a Conditional Probability Table (CPT) for variable C.

$$\Pr(c|a) + \Pr(\bar{c}|a) = 1 \text{ and } \Pr(c|\bar{a}) + \Pr(\bar{c}|\bar{a}) = 1.$$

Two of the probabilities in the above CPT are redundant and can be inferred from the other two. We only need 10 independent probabilities to completely specify the CPTs for this DAG.

Parameterizing the Independence Structure

Definition

A Bayesian network for variables $\mathbf{Z}$ is a pair (G, Θ), where

- G is a directed acyclic graph over variables $\mathbf{Z}$, called the network structure.
- Θ is a set of conditional probability tables (CPTs), one for each variable in $\mathbf{Z}$, called the network parametrization.

- $\Theta_{X|\mathbf{U}}$: the CPT for variable X and its parents $\mathbf{U}$.
- $X\mathbf{U}$: a network family.
- $\theta_{x|\mathbf{u}}$: the value assigned by CPT $\Theta_{X|\mathbf{U}}$ to the conditional probability $\Pr(x|\mathbf{u})$. Called a network parameter.

We must have $\sum_x \theta_{x|\mathbf{u}} = 1$ for every parent instantiation $\mathbf{u}$.

Parameterizing the Independence Structure

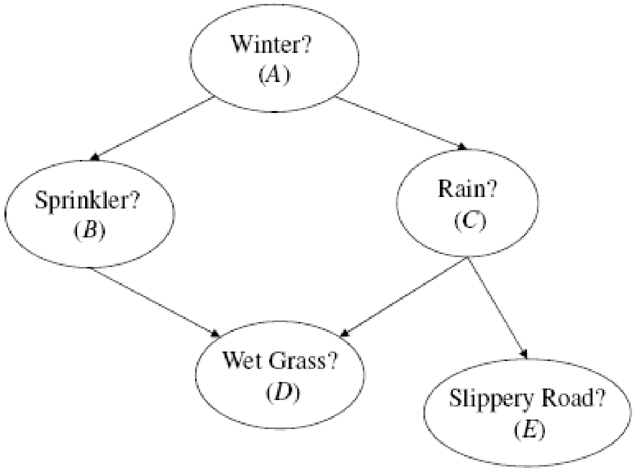

A	B	$\Theta_{B\|A}$
true	true	.2
true	false	.8
false	true	.75
false	false	.25

A	C	$\Theta_{C\|A}$
true	true	.8
true	false	.2
false	true	.1
false	false	.9

A	Θ_A
true	.6
false	.4

B	C	D	$\Theta_{D\|B,C}$
true	true	true	.95
true	true	false	.05
true	false	true	.9
true	false	false	.1
false	true	true	.8
false	true	false	.2
false	false	true	0
false	false	false	1

C	E	$\Theta_{E\|C}$
true	true	.7
true	false	.3
false	true	0
false	false	1

Use GeNie/Smile
To create this network

Parameterizing the Independence Structure

Chain rule for Bayesian networks

A Bayesian network is an implicit representation of a unique probability distribution $\Pr$ given by

$$\Pr(\mathbf{z}) \stackrel{def}{=} \prod_{\theta_{x|\mathbf{u}} \sim \mathbf{z}} \theta_{x|\mathbf{u}}.$$

The probability assigned to a network instantiation $\mathbf{z}$ is simply the product of all network parameters that are compatible with $\mathbf{z}$.

Parameterizing the Independence Structure

Chain rule for Bayesian networks

A Bayesian network is an implicit representation of a unique probability distribution $\Pr$ given by

$$\Pr(\mathbf{z}) \stackrel{def}{=} \prod_{\theta_{x|\mathbf{u}} \sim \mathbf{z}} \theta_{x|\mathbf{u}}.$$

The probability assigned to a network instantiation $\mathbf{z}$ is simply the product of all network parameters that are compatible with $\mathbf{z}$.

Parameterizing the Independence Structure

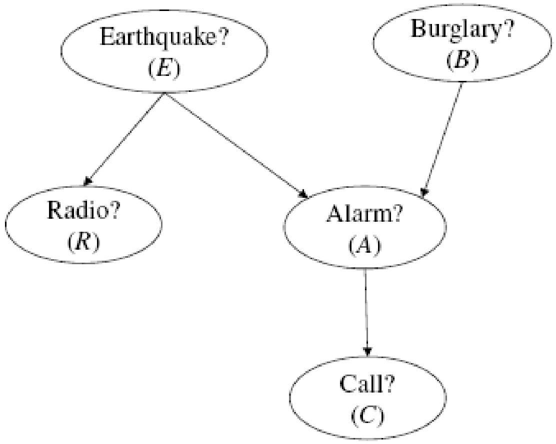

Example

$$
\begin{aligned}
&\Pr(a, b, \bar{c}, d, \bar{e}) \\
&\quad = \theta_a \; \theta_{b|a} \; \theta_{\bar{c}|a} \; \theta_{d|b,\bar{c}} \; \theta_{\bar{e}|\bar{c}} \\
&\quad = (.6)(.2)(.2)(.9)(1) \\
&\quad = .0216
\end{aligned}
$$

Example

$$
\begin{aligned}
&\Pr(\bar{a}, \bar{b}, \bar{c}, \bar{d}, \bar{e}) \\
&\quad = \theta_{\bar{a}} \; \theta_{\bar{b}|\bar{a}} \; \theta_{\bar{c}|\bar{a}} \; \theta_{\bar{d}|\bar{b},\bar{c}} \; \theta_{\bar{e}|\bar{c}} \\
&\quad = (.4)(.25)(.9)(1)(1) \\
&\quad = .09
\end{aligned}
$$

Parameterizing the Independence Structure

- The CPT $\Theta_{X|\mathbf{U}}$ is exponential in the number of parents $\mathbf{U}$.
- If every variable can take up to d values, and has at most k parents, the size of any CPT is bounded by $O(d^{k+1})$.
- If we have n network variables, the total number of Bayesian network parameters is bounded by $O(n \cdot d^{k+1})$.
- This number is quite reasonable as long as the number of parents per variable is relatively small.

Bayesian Networks: Representation

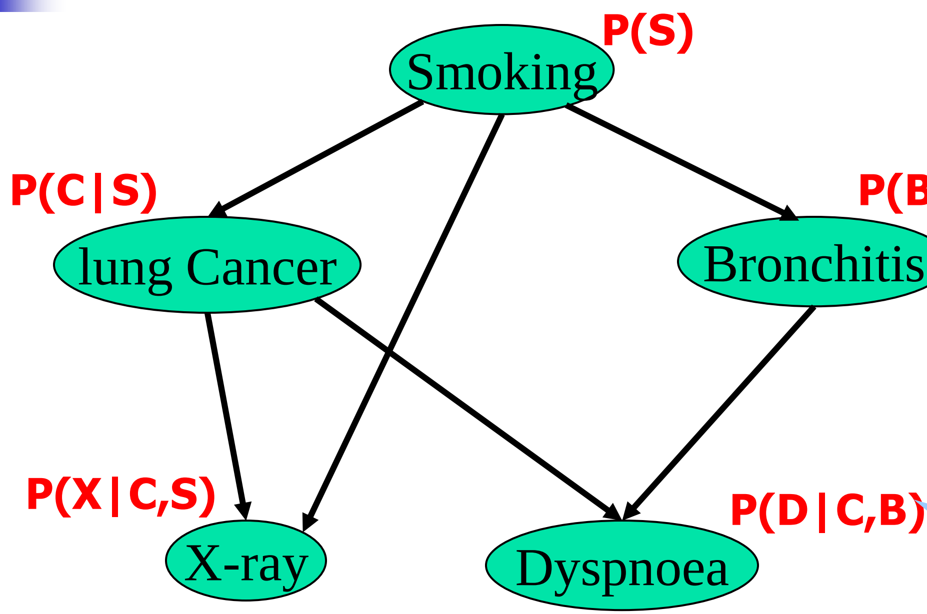

BN = (G, Θ)

CPD:

C	B	D=0	D=1
0	0	0.1	0.9
0	1	0.7	0.3
1	0	0.8	0.2
1	1	0.9	0.1

$$P(S, C, B, X, D) = P(S)\, P(C|S)\, P(B|S)\, P(X|C,S)\, P(D|C,B)$$

Conditional Independencies ➡ Efficient Representation

Outline

- Bayesian Networks, DAGS, Markov(G)
- Graphoids axioms for Conditional Independence
- D-separation: Inferring CIs in graphs

Properties of Probabilistic Independence

This independence follows from the Markov assumption

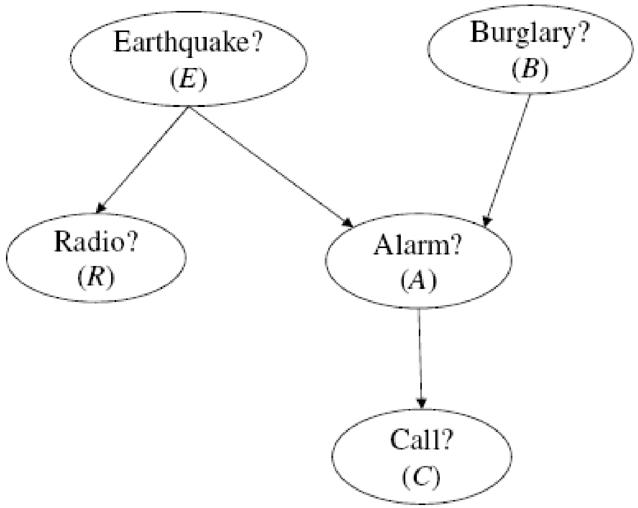

The distribution $\Pr$ specified by a Bayesian network (G, Θ) is guaranteed to satisfy every independence assumption in $\mathrm{Markov}(G)$.

These, however, are not the only independencies satisfied by the distribution $\Pr$.

R and C are independent given A

Properties of Probabilistic Independence
(Pearl ch 3)

THEOREM 1: Let X, Y, and Z be three disjoint subsets of variables from U. If $I(X, Z, Y)$ stands for the relation "X is independent of Y, given Z" in some probabilistic model P, then I must satisfy the following four independent conditions:

- Symmetry:
  - **I(X,Z,Y) → I(Y,Z,X)**
- Decomposition:
  - **I(X,Z,YW)→ I(X,Z,Y) and I(X,Z,W)**
- Weak union:
  - **I(X,Z,YW)→I(X,ZW,Y)**
- Contraction:
  - **I(X,Z,Y) and I(X,ZY,W)→I(X,Z,YW)**
- Intersection:
  - **I(X,ZY,W) and I(X,ZW,Y) → I(X,Z,YW)**

Symmetry

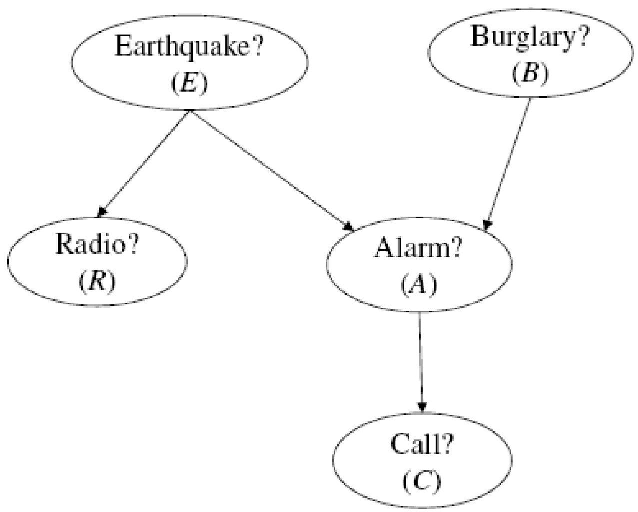

$$I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) \text{ iff } I_{\mathrm{Pr}}(\mathbf{Y}, \mathbf{Z}, \mathbf{X})$$

If learning **y** does not influence our belief in **x**, then learning **x** does not influence our belief in **y** either.

Example

From the independencies declared by $\mathrm{Markov}(G)$, we know that $I_{\mathrm{Pr}}(A, \{B, E\}, R)$. Using Symmetry, we can then conclude that $I_{\mathrm{Pr}}(R, \{B, E\}, A)$, which is not part of the independencies declared by $\mathrm{Markov}(G)$.

Decomposition

If some information is irrelevant, then any part of it is also irrelevant.

$$I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W}) \text{ only if } I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) \text{ and } I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z}, \mathbf{W}).$$

If learning **yw** does not influence our belief in **x**, then learning **y** alone, or learning **w** alone, will not influence our belief in **x** either.

Pearl's language:
If two pieces of information are irrelevant to X then each one is irrelevant to X

Decomposition

The opposite of Decomposition, called Composition:

$$I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) \text{ and } I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{W}) \cancel{\text{ only if }} \Longrightarrow I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W})$$

does not hold in general.

Two pieces of information may each be irrelevant on their own, yet their combination may be relevant.

Example: Two coins (C1,C2,) and a bell (B)

Decomposition

More generally...

Decomposition allows us to state the following:

$$I_{\mathrm{Pr}}(X, \mathrm{Parents}(X), \mathbf{W}) \quad \text{for every } \mathbf{W} \subseteq \mathrm{Non\_Descendants}(X).$$

Every variable X is conditionally independent of any subset of its non-descendants given its parents.

This is a strengthening of the independence statements declared by $\mathrm{Markov}(G)$, which is a special case when $\mathbf{W}$ contains all non-descendants of X.

Decomposition

Decomposition proves the chain rule for Bayesian networks.

By the chain rule of probability calculus:

$$\Pr(r, c, a, e, b) = \Pr(r|c, a, e, b)\Pr(c|a, e, b)\Pr(a|e, b)\Pr(e|b)\Pr(b).$$

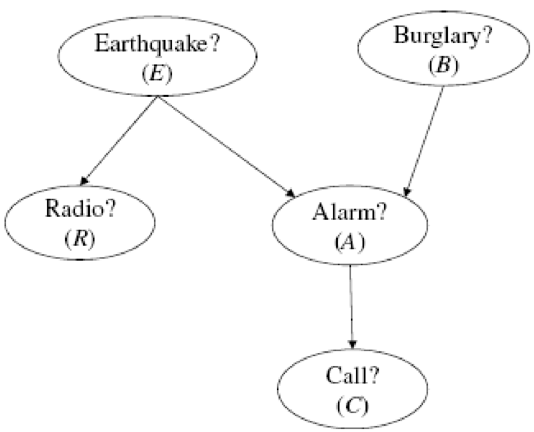

By Decomposition:

$$\begin{aligned}
\Pr(r|c, a, e, b) &= \Pr(r|e) \\
\Pr(c|a, e, b) &= \Pr(c|a) \\
\Pr(e|b) &= \Pr(e).
\end{aligned}$$

This leads to the chain rule of Bayesian networks:

$$\begin{aligned}
\Pr(r, c, a, e, b) &= \Pr(r|e)\Pr(c|a)\Pr(a|e, b)\Pr(e)\Pr(b) \\
&= \theta_{r|e}\ \theta_{c|a}\ \theta_{a|e,b}\ \theta_e\ \theta_b.
\end{aligned}$$

Weak Union

$$I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W}) \ \cancel{\text{only if}} \Longrightarrow \ I_{\Pr}(\mathbf{X}, \mathbf{Z} \cup \mathbf{Y}, \mathbf{W})$$

If the information **yw** is not relevant to our belief in **x**, then the partial information **y** will not make the rest of the information, **w**, relevant.

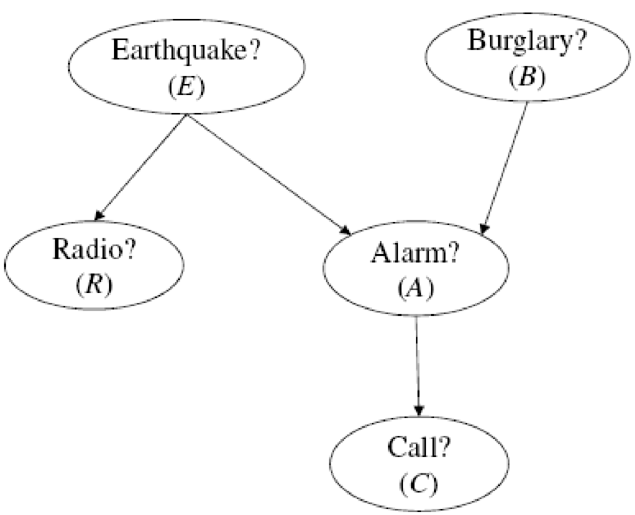

$I(C, A, \{B, E, R\})$ is part of $\text{Markov}(G)$. By Weak Union: $I_{\Pr}(C, \{A, B, E\}, R)$, which is not part of the independencies declared by $\text{Markov}(G)$.

Contraction

$$I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) \text{ and } I_{\Pr}(\mathbf{X}, \mathbf{Z} \cup \color{red}{\mathbf{Y}}, \mathbf{W}) \ \cancel{\text{only if}} \Longrightarrow I_{\Pr}(\mathbf{X}, \mathbf{Z}, \color{red}{\mathbf{Y}} \cup \mathbf{W})$$

If after learning the irrelevant information **y**, the information **w** is found to be irrelevant to our belief in **x**, then the combined information **yw** must have been irrelevant from the beginning.

Compare Contraction with Composition:

$$I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y}) \text{ and } I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{W}) \ \cancel{\text{only if}} \Longrightarrow I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W})$$

One can view Contraction as a weaker version of Composition. Recall that Composition does not hold for probability distributions.

Strictly Positive Distributions

When there are no constraints

Definition

A strictly positive distribution assign a non-zero probability to every consistent event.

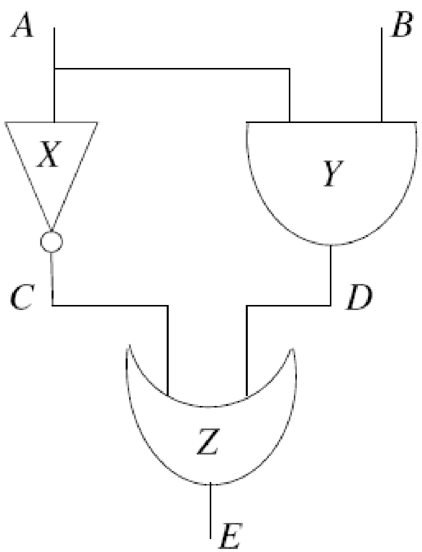

Example

A strictly positive distribution cannot represent the behavior of Inverter X as it will have to assign the probability zero to the event $A = \text{true}$, $C = \text{true}$.

A strictly positive distribution cannot capture logical constraints.

Intersection

Holds only for strictly positive distributions

$I_{\Pr}(\mathbf{X}, \mathbf{Z} \cup \mathbf{W}, \mathbf{Y})$ and $I_{\Pr}(\mathbf{X}, \mathbf{Z} \cup \mathbf{Y}, \mathbf{W})$ only if $I_{\Pr}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W})$

If information **w** is irrelevant given **y**, and **y** is irrelevant given **w**, then combined information **yw** is irrelevant to start with.

Intersection

Holds only for strictly positive distributions

$I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z} \cup \mathbf{W}, \mathbf{Y})$ and $I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z} \cup \mathbf{Y}, \mathbf{W})$ only if $I_{\mathrm{Pr}}(\mathbf{X}, \mathbf{Z}, \mathbf{Y} \cup \mathbf{W})$

If information **w** is irrelevant given **y**, and **y** is irrelevant given **w**, then combined information **yw** is irrelevant to start with.

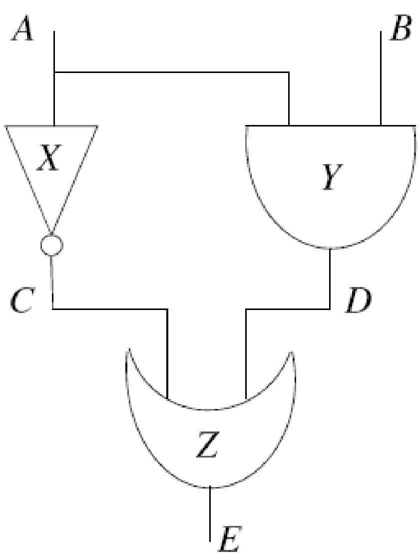

- If we know the input A of inverter X, its output C becomes irrelevant to our belief in the circuit output E.
- If we know the output C of inverter X, its input A becomes irrelevant to this belief.
- Yet, variables A and C are not irrelevant to our belief in the circuit output E.

Properties of Probabilistic independence

THEOREM 1: Let X, Y, and Z be three disjoint subsets of variables from U. If $I(X, Z, Y)$ stands for the relation "X is independent of Y, given Z" in some probabilistic model P, then I must satisfy the following four independent conditions:

- Symmetry:
  - **I(X,Z,Y) → I(Y,Z,X)**
- Decomposition:
  - **I(X,Z,YW)→ I(X,Z,Y) and I(X,Z,W)**
- Weak union:
  - **I(X,Z,YW)→I(X,ZW,Y)**
- Contraction:
  - **I(X,Z,Y) and I(X,ZY,W)→I(X,Z,YW)**
- Intersection:
  - **I(X,ZY,W) and I(X,ZW,Y) → I(X,Z,YW)**

Graphoid axioms:
Symmetry, decomposition
Weak union and contraction

Positive graphoid:
+intersection

In Pearl: the 5 axioms
are called Graphids,
the 4, semi-graphois