

Probabilistic Graphical Models: Representation and Inference

Aaron C. Courville
Université de Montréal

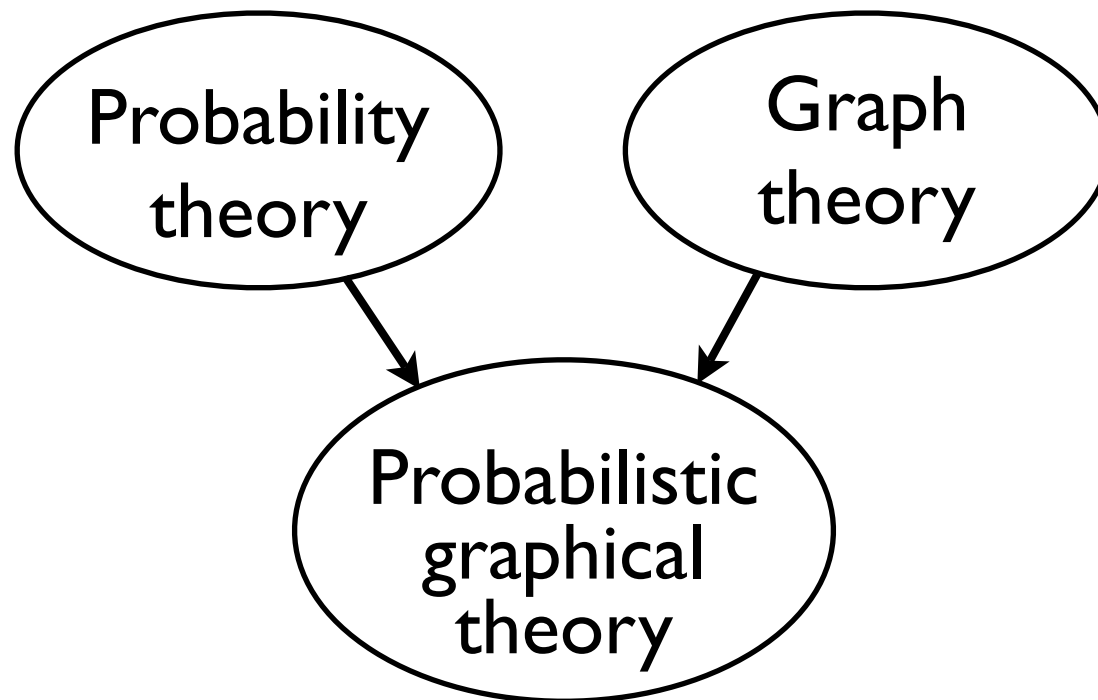
Note: Material for the slides is taken directly from a presentation prepared by Andrew Moore

Overview

- Graphs and probabilities
 - Conditional independence
 - Directed graphs
 - Examples
- Inference in a Directed Graphical Model
 - Belief propagation
- Learning
 - Maximum likelihood I: fully observable case
 - Maximum likelihood II: Expectation Maximum
- Approximate Inference
 - Variational Methods
 - Monte Carlo Methods

Probabilistic Graphical Models

- Graphs endowed with a probability distribution
 - nodes represent random variables and the edges encode conditional independence assumptions
- Graphical model express sets of conditional independence via graph structure (and conditional independence is useful)
- Graph structure plus associated parameters define joint probability distribution of the set of nodes/variables



Graphical Models

- Graphical models come in two main flavors:
 1. Directed graphical models (a.k.a Bayes Net, Belief Networks):
 - consists of a set of nodes with arrows (directed edges) between some of the nodes
 - arrows encode factorized conditional probability distributions
 2. Undirected graphical models (a.k.a Markov random fields):
 - consists of a set of nodes with undirected edges between some of the nodes
 - edges (or more accurately the lack of edges) encode conditional independence.
- Today, we will focus exclusive on directed graphs.

Probability Review: Marginal Independence

Definition: X is marginally independent of Y if for all (i, j)

$$\begin{aligned}P(X = x_i, Y = y_j) &= P(X = x_i)P(Y = y_j) \\ P(X, Y) &= P(X)P(Y)\end{aligned}$$

This is equivalent to the expression:

$$P(X | Y) = P(X) \quad P(Y | X) = P(Y)$$

Why? Recall from the probability product rule

$$P(X, Y) = P(X | Y)P(Y) = P(Y | X)P(X) = P(X)P(Y)$$

Probability Review:

Conditional Independence

Definition: X is conditionally independent of Y given Z if the probability distribution governing X is independent of the value of Y , given the value of Z : for all (i, j, k)

$$\begin{aligned}P(X = x_i, Y = y_j \mid Z = z_k) &= P(X = x_i \mid Z = z_k)P(Y = y_j \mid Z = z_k) \\P(X, Y \mid Z) &= P(X \mid Z)P(Y \mid Z)\end{aligned}$$

Or equivalently (by the product rule):

$$P(X \mid Y, Z) = P(X \mid Z) \quad P(Y \mid X, Z) = P(Y \mid Z)$$

Why? Recall from the probability product rule

$$P(X, Y, Z) = P(X \mid Y, Z)P(Y \mid Z)P(Z) = P(X \mid Z)P(Y \mid Z)P(Z)$$

Example: $P(\textit{Thunder} \mid \textit{Rain}, \textit{Lightning}) = P(\textit{Thunder} \mid \textit{Lightning})$

Directed Graphical Models

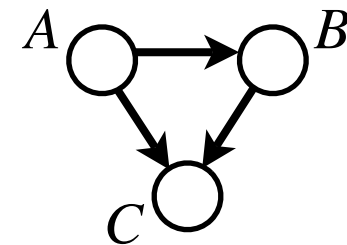
- Set of nodes with arrows (directed edges) between some of the nodes
 - Nodes represent random variables
 - Arrows encode factorized conditional probability distributions

- Consider an arbitrary joint distribution:

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

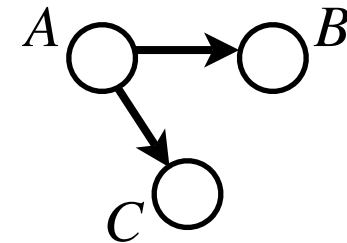
- By application of the product rule:

$$\begin{aligned} P(A, B, C) &= P(A)P(B, C | A) \\ &= P(A)P(B | A)P(C | A, B) \end{aligned}$$

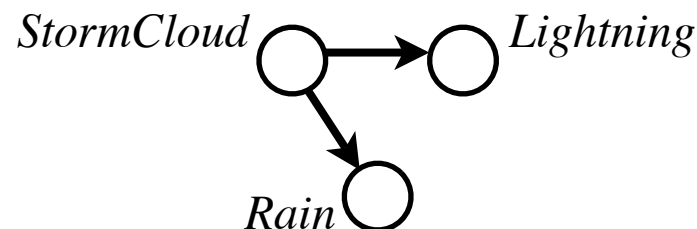


- What if C is conditionally independent from B :

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



E.g., $P(\text{Rain}, \text{Lightning} | \text{StormCloud}) = P(\text{Rain} | \text{StormCloud})P(\text{Lightning} | \text{StormCloud})$

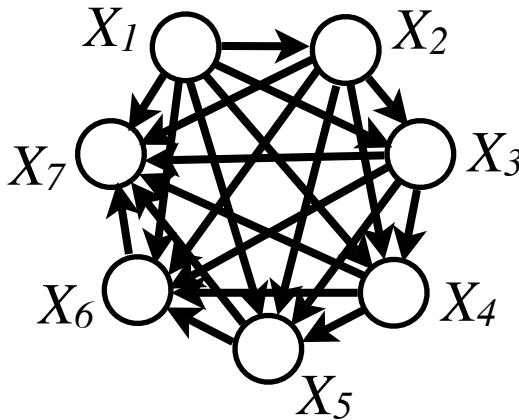


The General Case

- Consider the arbitrary joint distribution: $P(X_1, \dots, X_d)$
- By successive application of the product rule

$$P(X_1, \dots, X_d) = P(X_1)P(X_2 | X_1) \cdots P(X_d | X_1, \dots, X_{d-1})$$

- Can be represented by a graph in which each node has links from all lower-numbered nodes (i.e. *a fully connected graph*)



- *Directed acyclic graph*: NO DIRECTED CYCLES \Rightarrow can number nodes so that there are no links from higher numbered to lower numbered nodes.

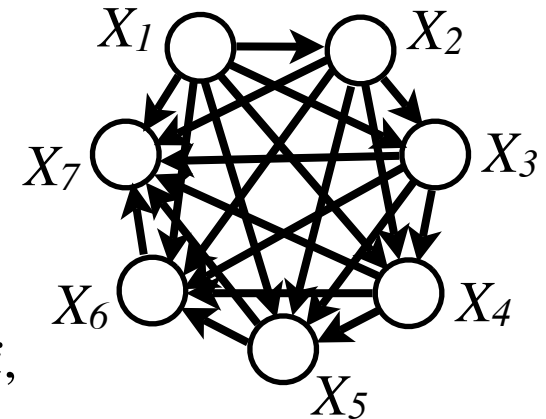
Directed Acyclic Graphs (DAGs)

- General factorization:

$$P(X_1, \dots, X_d) = \prod_{i=1}^d P(X_i \mid Pa_i)$$

where Pa_i denotes the immediate parents of node i ,

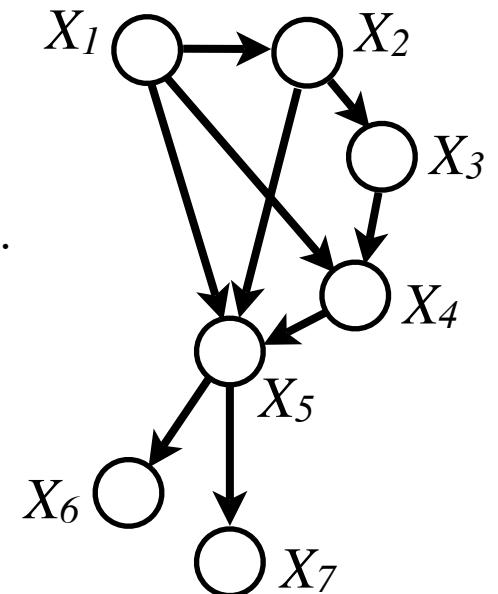
- i.e. all the nodes with arrows leading to node i .



- Node X_i is conditionally independent of its non-descendants given its immediate parents (missing links imply conditional independence).

- Some terminology:

- Parents of node i = nodes with arrows **leading to** node i .
- Antecedents = parents, parents of parents, etc.
- Children of node i = nodes with arrows **leading from** node i .
- Descendants = children, children of children, etc.



Importance of Ordering

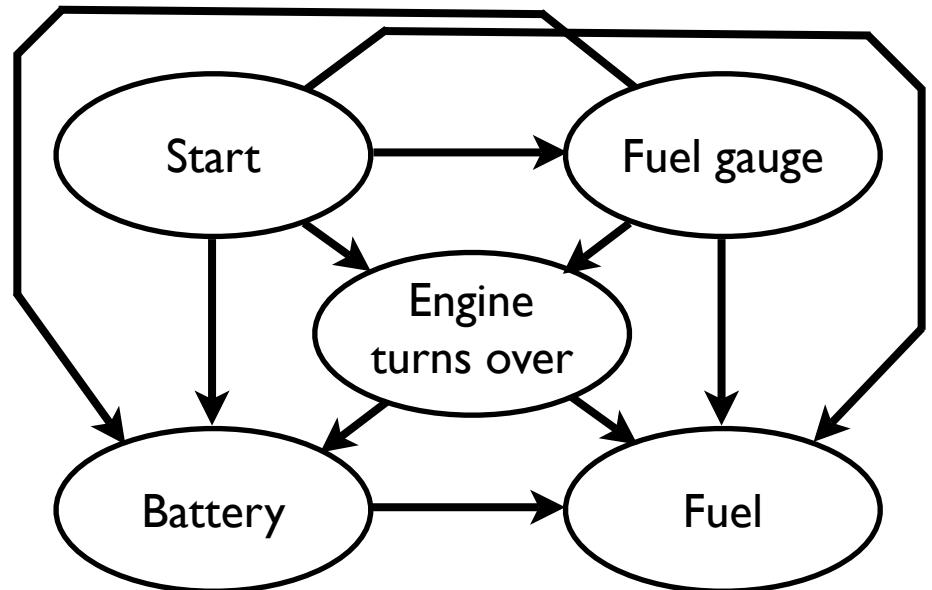
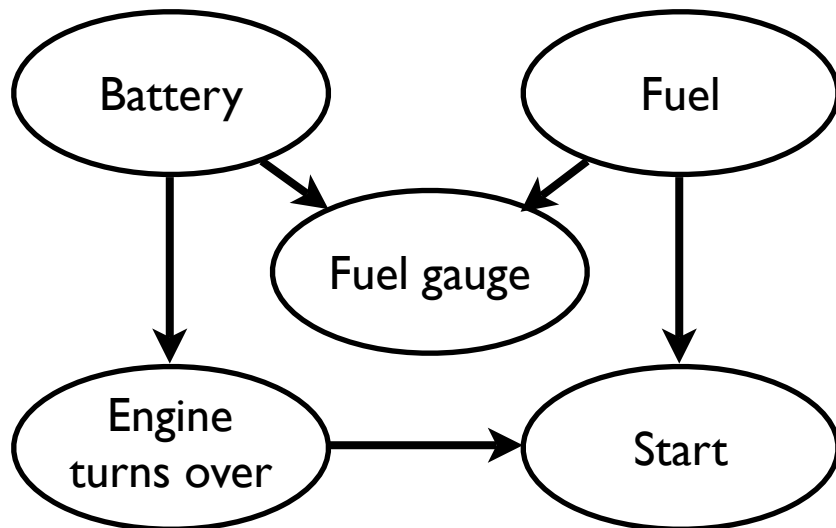
- Since *every* joint probability distribution can be decomposed as

$$P(X_1, \dots, X_d) = \prod_{i=1}^d P(X_i \mid Pa_i)$$

- then every joint probability can be represented as a graphical model
 - with arbitrary ordering over the variables
 - i.e. any decomposition will generate a valid graph

- However **not all decompositions are created equal.**

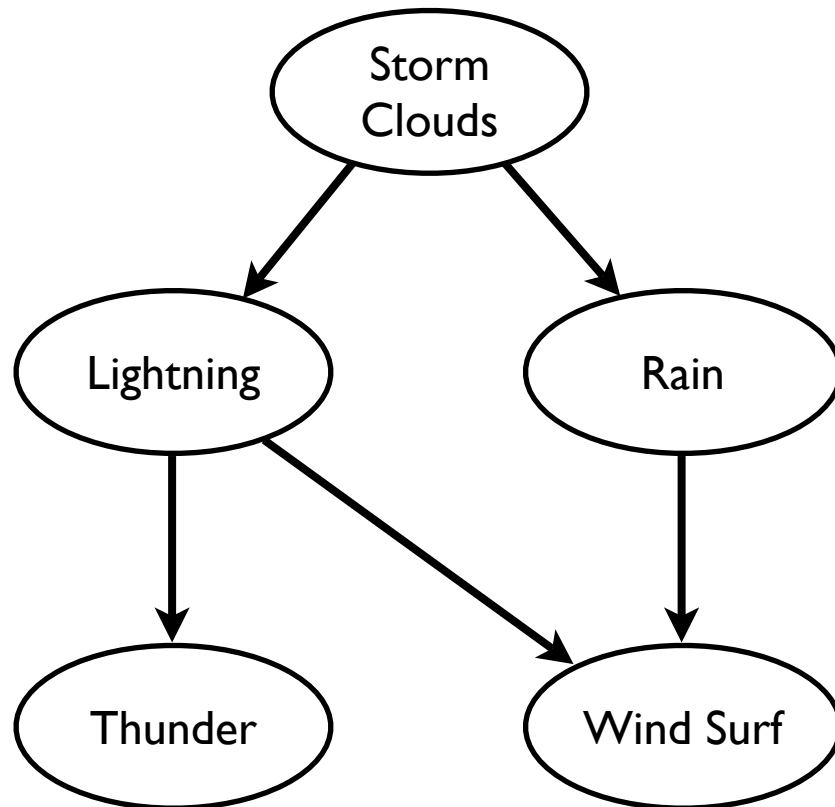
- An illustration:



DAG Example I: Bayes Net

A conditional probability distribution (CPD) is associated with each node.

- Defines $P(X_i | Pa_i)$.
- Specifies how the parents interact.



CPD for node: *Lightning*

Pa	$P(L SC)$	$P(\sim L SC)$
SC	0.7	0.3
$\sim SC$	0	1.0

CPD for node: *Wind Surf*

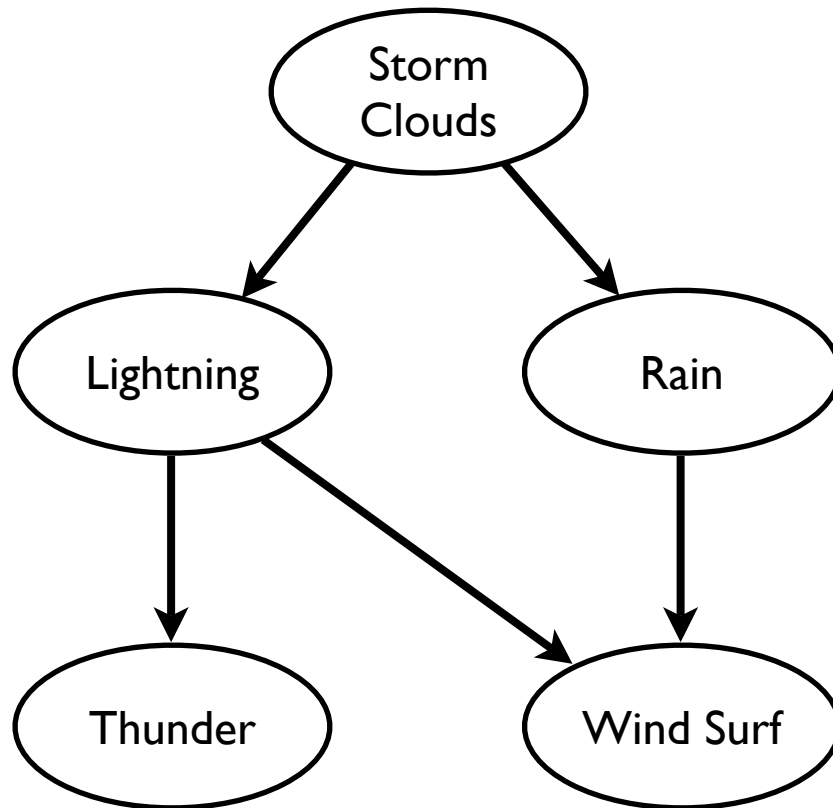
Pa	$P(WS Pa)$	$P(\sim WS Pa)$
L, R	0	1.0
L, $\sim R$	0	1.0
$\sim L$, R	0.2	0.8
$\sim L$, $\sim R$	0.9	0.1

$$\begin{aligned}
 P(SC, L, R, T, WS) &= \prod_{i=1}^d P(\text{node}_i | Pa_i) \\
 &= P(SC)P(L | SC)P(R | SC)P(T | L)P(WS | L, R)
 \end{aligned}$$

Example I: Bayes Net (cont.)

Let's count the parameters

- In the full joint distribution: $2^5 - 1 = 31$
- Taking advantage of conditional independences: 11



CPD for node: *Lightning*

Pa	$P(L SC)$	$P(\sim L SC)$
SC	0.7	0.3
$\sim SC$	0	1.0

CPD for node: *Wind Surf*

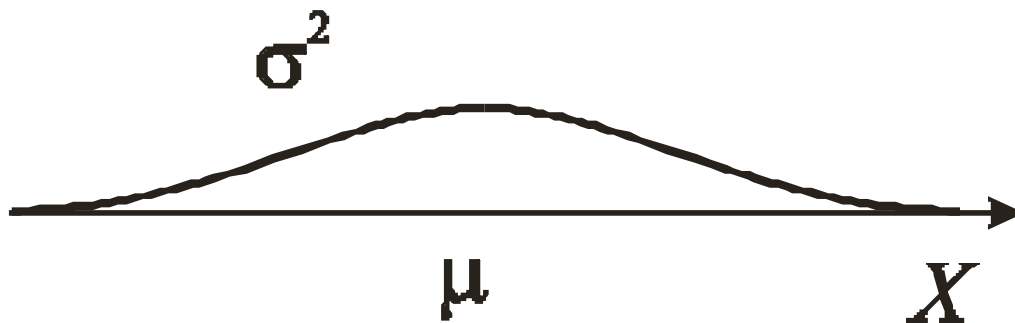
Pa	$P(WS Pa)$	$P(\sim WS Pa)$
L, R	0	1.0
L, $\sim R$	0	1.0
$\sim L$, R	0.2	0.8
$\sim L$, $\sim R$	0.9	0.1

$$P(SC, L, R, T, WS) = P(SC)P(L | SC)P(R | SC)P(T | L)P(WS | L, R)$$

Example: Univariate Gaussian

- Consider a univariate Gaussian random variable as a (simple) graphical model.

$$p(X = x) = \mathcal{N}(x \mid \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right)$$



Graphical model:

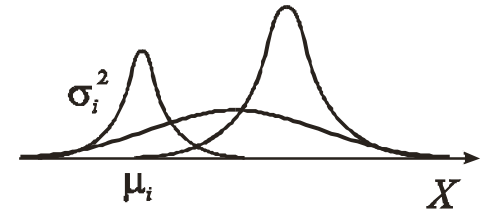
$X \circ$

Example II: Mixture of Gaussians

- Now let's consider a random variable distributed according to a mixture of Gaussians.
- **Conditional distributions:**

$$P(I = i) = w_i$$

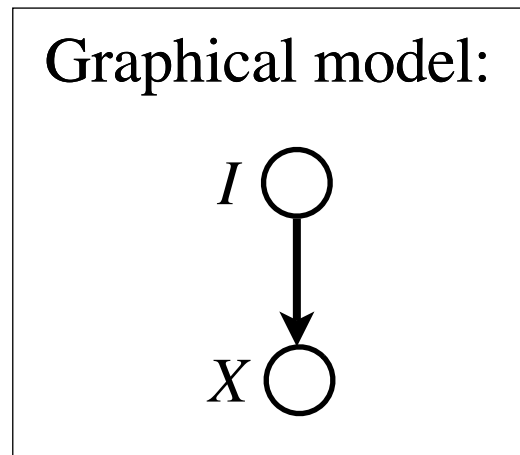
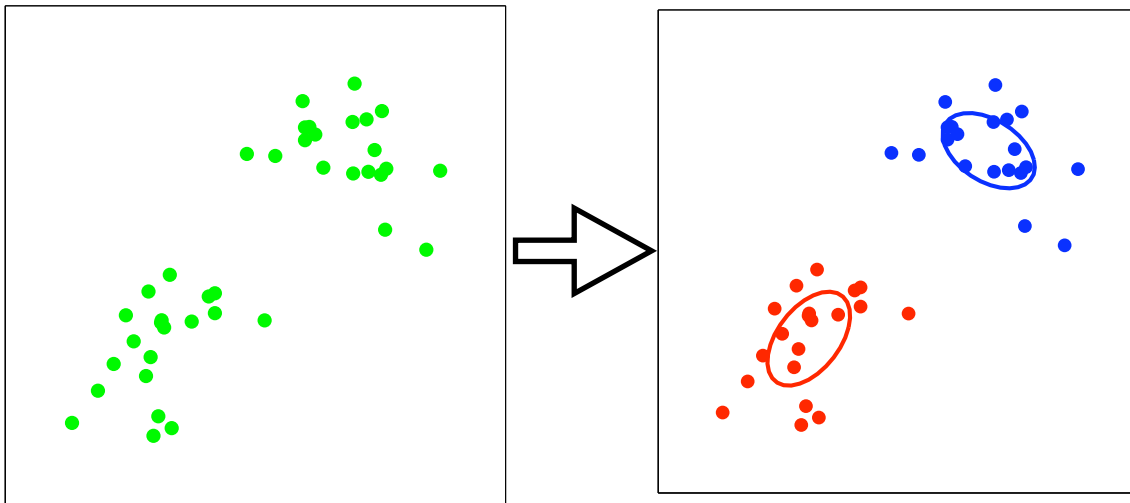
$$p(X = x | I = i) = \mathcal{N}(x | \mu_i, \sigma_i^2) = \frac{1}{\sigma_i \sqrt{2\pi}} \exp\left(-\frac{(x - \mu_i)^2}{2\sigma_i^2}\right)$$



where I is an index over the Gaussian components in the mixture and the mixing proportion, w_i , is the marginal probability that X is generated by mixture component i .

- **Marginal distributions:**

$$p(X = x) = \sum_i p(X = x | I = i)P(I = i) = \sum_i w_i \mathcal{N}(x | \mu_i, \sigma_i^2)$$



Example III: Naïve Bayes

- Consider the classification problem:

X Y

Training data:

Sky	Temp	Humid	Wind	Water	Forecst	EnjoySpt
Sunny	Warm	Normal	Strong	Warm	Same	Yes
Sunny	Warm	High	Strong	Warm	Same	Yes
Rainy	Cold	High	Strong	Warm	Change	No
Sunny	Warm	High	Strong	Cool	Change	Yes

- For a new example X^{new} , We want to determine $P(Y | X^{new})$ so that we can classify Y^{new} as $\operatorname{argmax}_j P(Y = j | X^{new})$.
- One way to do this is to use Bayes Rule:

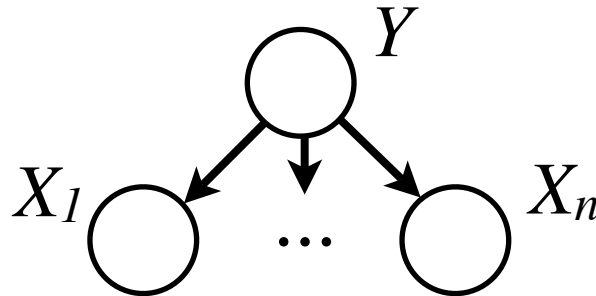
$$P(Y | X) = \frac{P(X | Y)P(Y)}{P(X)}$$

Example III: Naïve Bayes

- Using Bayes rule, to compute $P(Y | X)$ we need to represent $P(X | Y)$ and $P(Y)$.
 - For even marginally large X , the full joint $P(X_1, X_2, \dots, X_n | Y)$ is impractical
 - We would require a parameter for each unique combination of X_1, \dots, X_n .
- Naïve Bayes assumption:

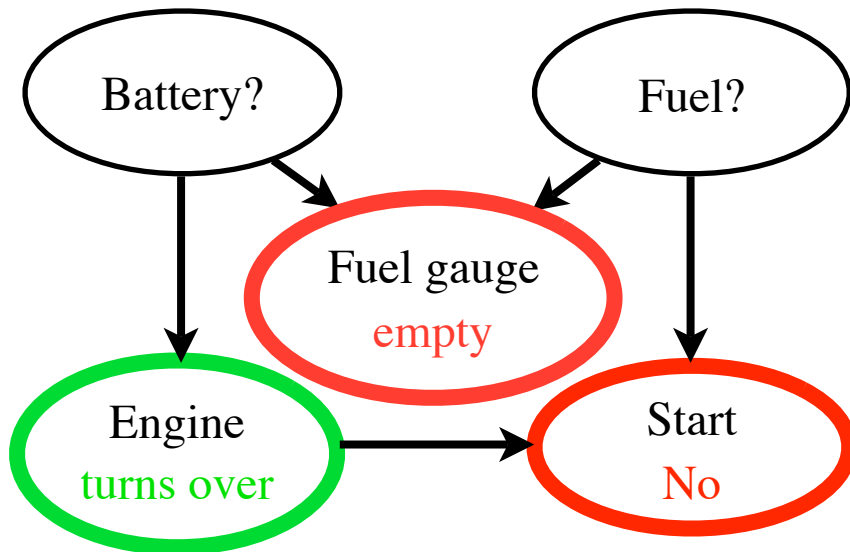
$$P(X | Y) = P(X_1, \dots, X_n | Y) = \prod_i P(X_i | Y)$$

- X_i and X_j ($i \neq j$) are conditionally independent given Y .



Conditioning on Evidence

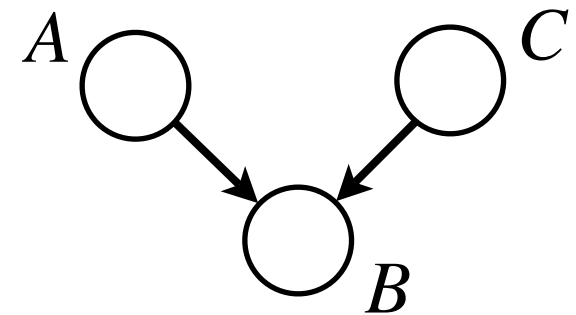
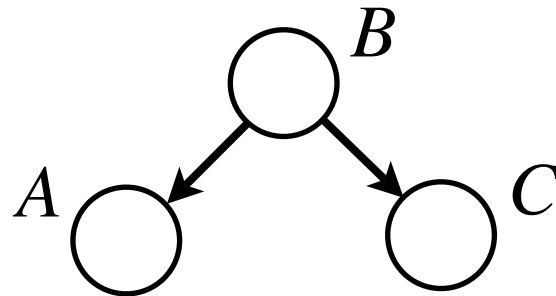
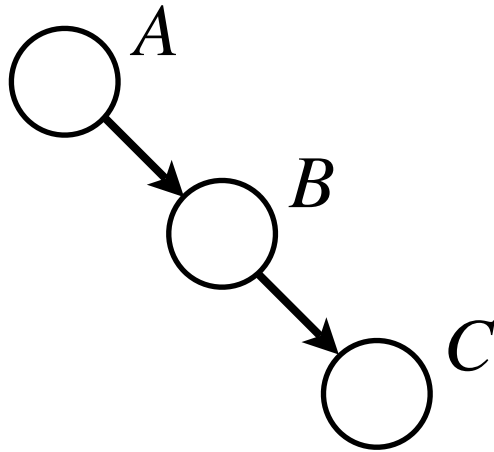
- Much of the point of Bayes Nets (and Graphical models, more generally) is to use data or observed variables to *infer* the values of hidden (or latent) variables.
- We can condition on the observed variables and compute conditional distributions over the latent variables.



Note: hidden variables may have a specific interpretation (such as Battery and Fuel), or may be introduced to permit a richer class of distributions (as in neural network hidden units).

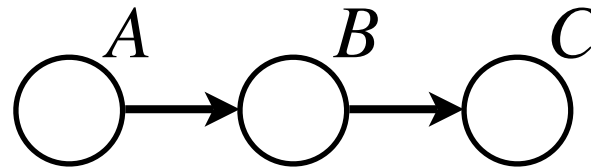
Markov Properties

- We have said that the conditional independence properties are represented directly from the graph.
- The pattern of conditional independence turn out to be extremely important for determining how to calculate the conditional distributions over variables of interest.
- We will now consider three simple independence structures that form the basic building blocks of all independence relationships representable by directed graphical models.



Markov Properties I

- Consider the joint distribution, $P(A,B,C)$ specified by the graph:

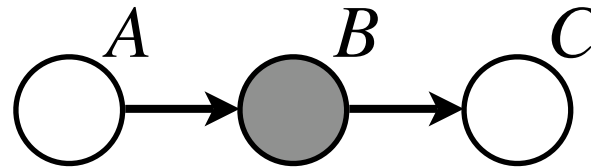


- Note the missing edge from A to C .
- Node B is “head-to-tail” with respect to the path A - B - C .
- Joint distribution:

$$\begin{aligned} P(A, B, C) &= P(A)P(B | A)P(C | A, B) \\ &= P(A)P(B | A)P(C | B) \end{aligned}$$

Markov Properties I (cont.)

- Suppose we condition on the node B :



- A and C are conditionally independent, given B : $A \perp\!\!\!\perp C \mid B$

$$\begin{aligned} P(A, C \mid B) &= \frac{P(A, B, C)}{P(B)} \\ &= \frac{P(A)P(B \mid A)P(C \mid B)}{P(B)} && \text{[By definition of joint]} \\ &= \frac{P(A)P(B \mid A)}{P(B)} P(C \mid B) \\ &= P(A \mid B)P(C \mid B) && \text{[By Bayes rule]} \end{aligned}$$

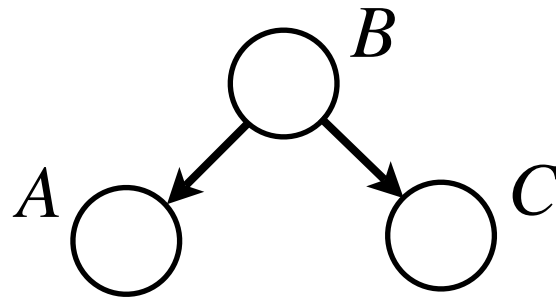
- Note that if B is not observed: $A \not\perp\!\!\!\perp C \mid \emptyset$

$$P(A, C) = \sum_B P(A, C \mid B)P(B) = \sum_B P(A \mid B)P(C \mid B)P(B) \neq P(A)P(C)$$

- Informally: observing B “blocks the path” from A to C .

Markov Properties II

- Consider the graph:

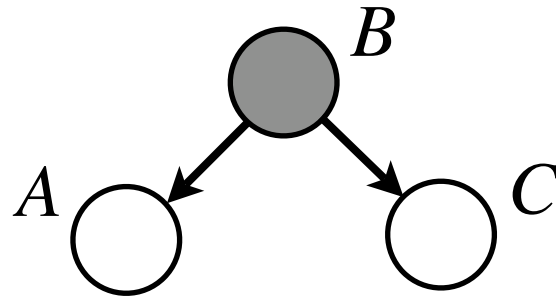


- Again, note missing edge from A to C .
- Node B is “tail-to-tail” with respect to the path A - B - C .
- Joint distribution:

$$\begin{aligned} P(A, B, C) &= P(B)P(A | B)P(C | A, B) \\ &= P(B)P(A | B)P(C | B) \end{aligned}$$

Markov Properties II (cont.)

- Again, let's condition on node B :



- A and C are conditionally independent, given B: $A \perp\!\!\!\perp C \mid B$

$$P(A, C \mid B) = P(A \mid B)P(C \mid B)$$

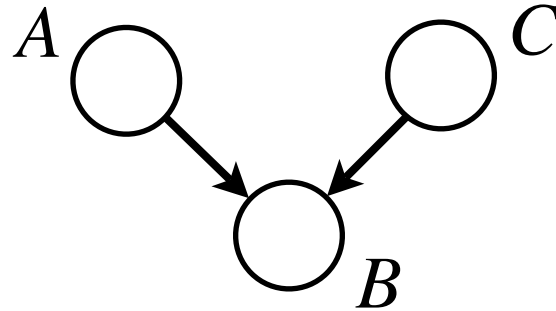
- Again, if B is not observed: $A \not\perp\!\!\!\perp C \mid \emptyset$

$$P(A, C) = \sum_B P(A, C \mid B)P(B) = \sum_B P(A \mid B)P(C \mid B)P(B) \neq P(A)P(C)$$

- Informally: observing B “blocks the path” from A to C .

Markov Properties III

- Consider the graph:



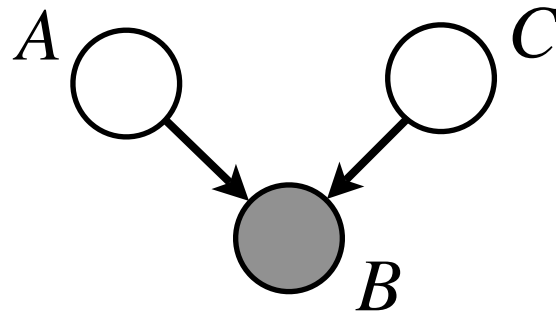
- Once again, note missing edge from A to C .
- Node B is “head-to-head” with respect to the path A - B - C .
- Joint distribution: $P(A, B, C) = P(A)P(C)P(B | A, C)$
- If B is not observed, we have: $A \perp\!\!\!\perp C | \emptyset$

$$\begin{aligned} P(A, C) &= \sum_B P(A)P(C)P(B | A, C) \\ &= P(A)P(C) \sum_B P(B | A, C) \\ &= P(A)P(C) \end{aligned}$$

Markov Properties III (cont.)

The Famous: “Explaining Away Effect”

- Suppose, once again, we condition on node B :



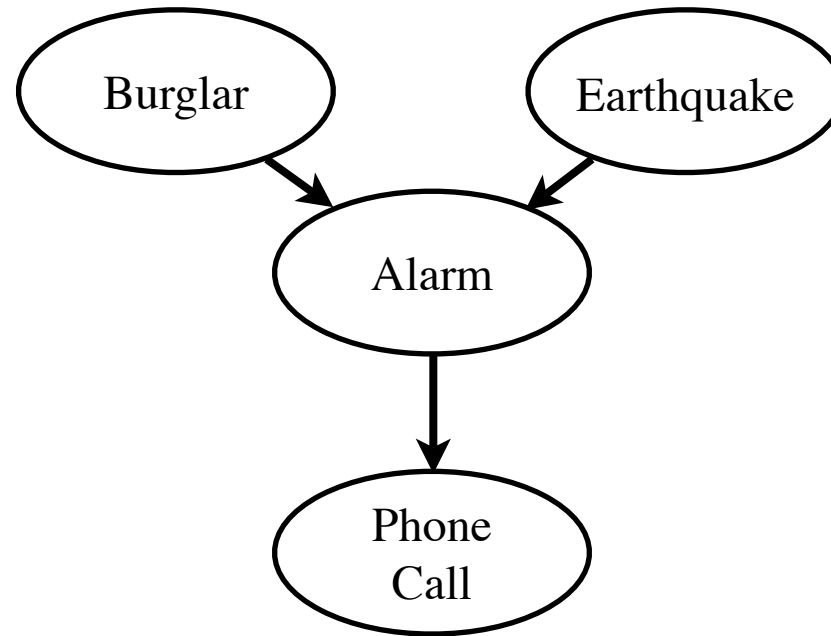
- A and C are *not* conditionally independent, given B : $A \not\perp C \mid B$

$$\begin{aligned} P(A, C \mid B) &= P(A \mid B, C)P(C \mid B) \\ &= P(C \mid B, A)P(A \mid B) \\ &\neq P(A \mid B)P(C \mid B) \end{aligned}$$

- Informally, an *unobserved* head-to-head node B “blocks the path” from A to C , but once B is observed the path is unblocked.
- **Important Note:** observation of *any descendent* of B also “unblocks the path.”

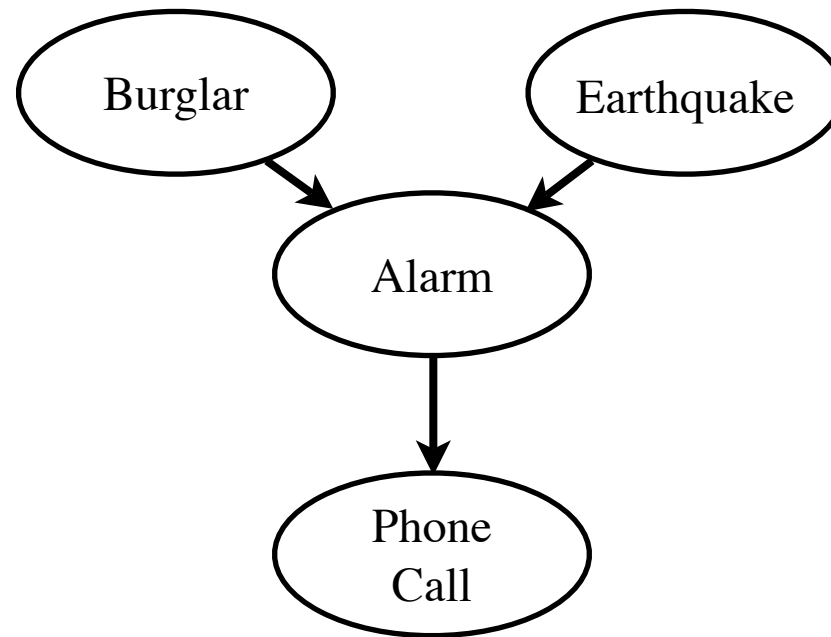
Explaining Away: An Illustration

- Consider the “Burglar Alarm” problem:



- Your house has a twitchy burglar alarm that is sometimes triggered by earthquakes.
- Earth doesn't care whether your house is currently being burgled
- While on vacation, one of your neighbours call and tells you that your home's burglar alarm is ringing!

The plot thickens...

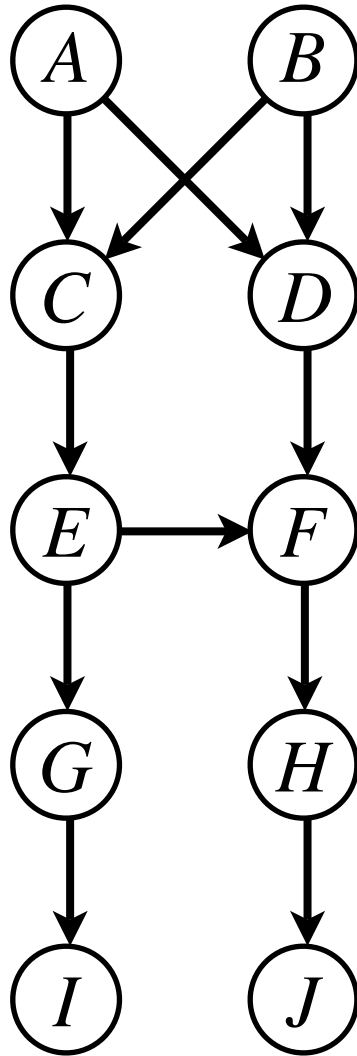


- But now suppose you learn that there was a small earthquake in your neighbourhood. Whew! Probably not a burglar after all.
- The Earthquake “explains away” the hypothetical burglar.
- So Burglar must not be conditionally independent of Earthquake given Alarm (or Phone Call).

d-Separation

- In a general DAG, consider three *groups* of nodes A , B , C .
- To determine whether A is conditional independent of C given B , consider all possible paths from any node in A to any node in C .
- Any such path is blocked if there is a node X which is head-to-tail or tail-to-tail with respect to the path and $X \in B$ or if the node is head-to-head and neither the node, *nor any of its descendants*, is in B .
- Note that a particular node may, for example, be head-to-head with respect to one particular path and head-to-tail with respect to a different path
- **Theorem** [Verma & Pearl, 1998]: If, given B , all possible paths from A to C are blocked (we say that B d-separates A and C), then A is conditionally independent of C , given B .
- Note: Variables may actually be independent even when not d-separated.

d-Separation Example



Conditional Independence

$C \perp\!\!\!\perp D ?$

$C \perp\!\!\!\perp D \mid A ?$

$C \perp\!\!\!\perp D \mid A, B ?$

$C \perp\!\!\!\perp D \mid A, B, J ?$

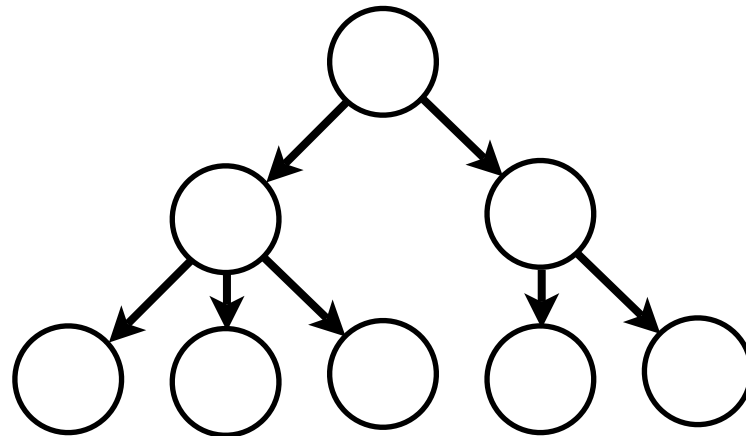
$C \perp\!\!\!\perp D \mid A, B, E, J ?$

Inference in a DAG

- Inference: calculate $P(X | Y)$ for some variable or sets of variables of interest X and observed variable or set of variables Y .
- If Z are the set of remaining variables in the graph:

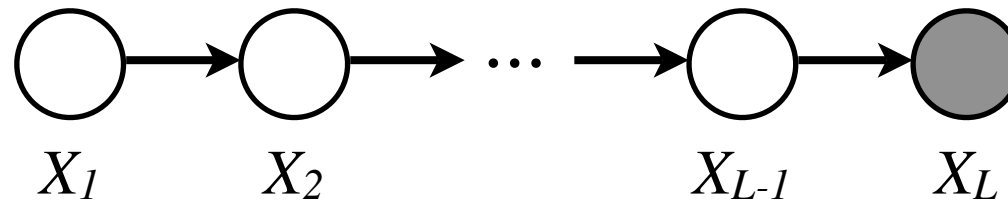
$$P(X | Y) = \sum_Z P(X | Z, Y)P(Z | Y)$$

- The bad news: exact inference in Bayes Nets (DAGs) is NP-hard!
- But exact inference is still tractable in some cases.
- Let's look at a special class of networks: *trees / forests*
 - Each node has at most one parent.



Example: Markov Chain

- Consider the graph: (a Markov Chain)



$$P(X_1, \dots, X_L) = P(X_1)P(X_2 | X_1) \dots P(X_L | X_{L-1})$$

- What if we want to find $P(X_1 | X_L)$?
- Direct evaluation gives:

$$P(X_1 | X_L) = \frac{\sum_{X_2} \dots \sum_{X_{L-1}} P(X_1 \dots X_L)}{\sum_{X_1} \sum_{X_2} \dots \sum_{X_{L-1}} P(X_1 \dots X_L)}$$

- For variables having M states, the denominator involves summing over M^{L-1} terms (exponential in the length of the chain).

Example: Markov Chain (cont.)

- Using the conditional independence structure we can reorder the summations in the denominator to give

$$\sum_{X_1} \sum_{X_2} \cdots \sum_{X_{L-1}} P(X_1 \dots X_L) = \sum_{X_{L-1}} P(X_L | X_{L-1}) \cdots \sum_{X_2} P(X_3 | X_2) \sum_{X_1} P(X_2 | X_1) P(X_1)$$

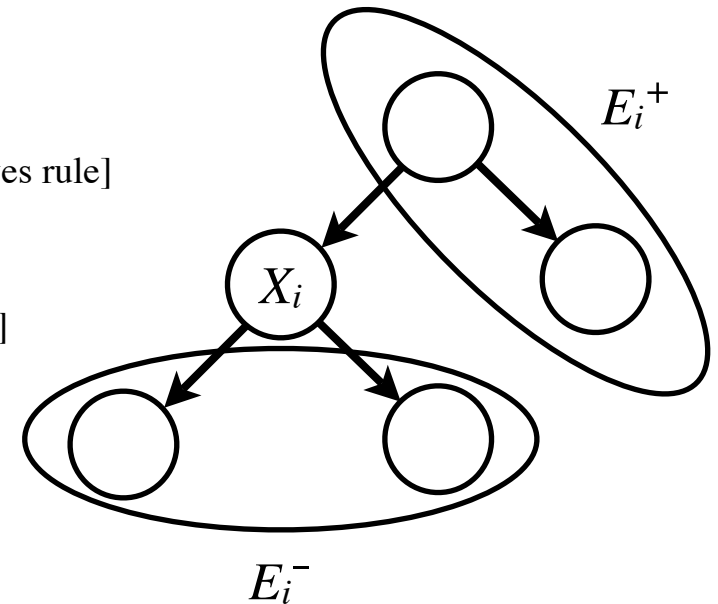
which involves on the order of LM^2 summations (linear in the length of the chain) - similarly for the numerator

- This is known as *variable elimination*, also can be interpreted as a sequence of messages being passed.

Belief Propagation: Decomposing Probabilities

- Suppose we want to compute $P(X_i | E)$ where E is some set of *evidence variables*.
 - Evidence variables are those that are observable, offering “evidence” with respect to the values of the latent variables.
- Let’s split E into two parts:
 - E_i^- is the part consisting of assignments to variables in the subtree rooted at X_i .
 - E_i^+ is the rest of it.

$$\begin{aligned}
 P(X_i | E) &= P(X_i | E_i^-, E_i^+) \\
 &= \frac{P(E_i^- | X_i, E_i^+) P(X_i | E_i^+)}{P(E_i^- | E_i^+)} && \text{[by Bayes rule]} \\
 &= \frac{P(E_i^- | X_i) P(X_i | E_i^+)}{P(E_i^- | E_i^+)} && \text{[by cond. ind.]} \\
 &= \alpha \pi(X_i) \lambda(X_i)
 \end{aligned}$$



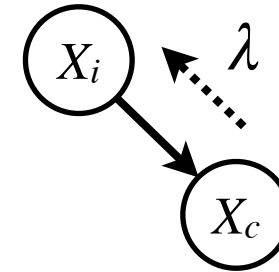
where $\pi(X_i) = P(X_i | E_i^+)$, $\lambda(X_i) = P(E_i^- | X_i)$
 and α is a constant independent of X_i .

Computing $\lambda(X_i)$ for leaves

- Starting at the leaves of the tree, recursively compute $\lambda(X_i) = P(E_i^- \mid X_i)$ for all X_i .
- If X_i is a leaf:
 - If X_i is in E (X_i is observed): $\lambda(X_i) = 1$ if X_i matches E , 0 otherwise.
 - If X_i is not in E : E_i^- is the null set, so $\lambda(X_i) = 1$ (constant).
- For simplicity, and without loss of generality, we will assume that *all* variables in E (the evidence set) are leaves in the tree.
 - How can we do this? In the case of a tree, conditioning on E d-separates the graph at that point.

Computing $\lambda(X_i)$ for non-leaves I

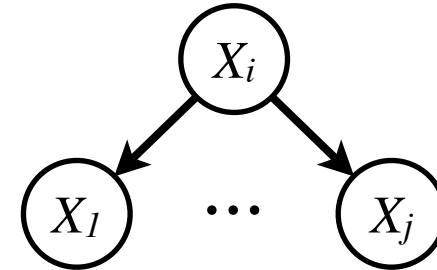
- Suppose X_i has one child, X_c .



$$\begin{aligned}\lambda(X_i) &= P(E_i^- | X_i) \\ &= \sum_j P(E_i^-, X_c = j | X_i) \\ &= \sum_j P(X_c = j | X_i) P(E_i^- | X_i, X_c = j) \\ &= \sum_j P(X_c = j | X_i) P(E_i^- | X_c = j) \\ &= \sum_j P(X_c = j | X_i) \lambda(X_c = j)\end{aligned}$$

Note: λ is a function of X (for discrete X , λ is a vector).

Computing $\lambda(X_i)$ for non-leaves II

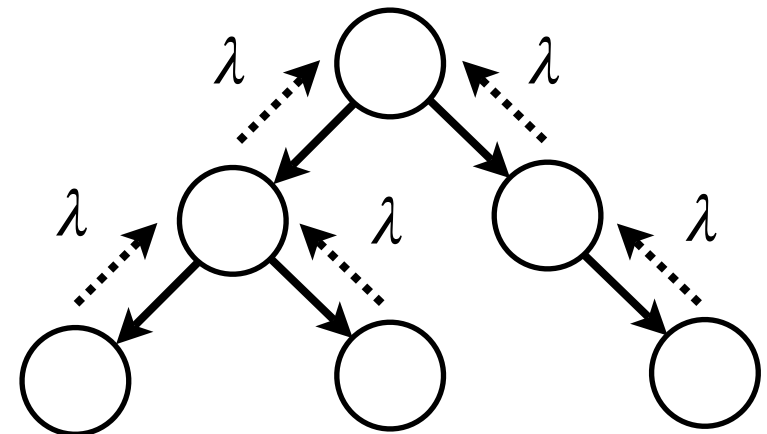


- Now suppose X_i has a set of children C .
- Since X_i d-separates each of its subtrees, the contribution of each subtree to $\lambda(X_i)$ is independent.

$$\lambda(X_i) = P(E_i^- | X_i) = \prod_{X_j \in C} \lambda_j(X_i) = \prod_{X_j \in C} \left[\sum_k P(X_j = k | X_i) \lambda(X_j) \right]$$

where $\lambda_j(X_i)$ is the contribution to $\lambda(X_i)$ of the evidence lying in the subtree rooted at one of X_i 's children X_j .

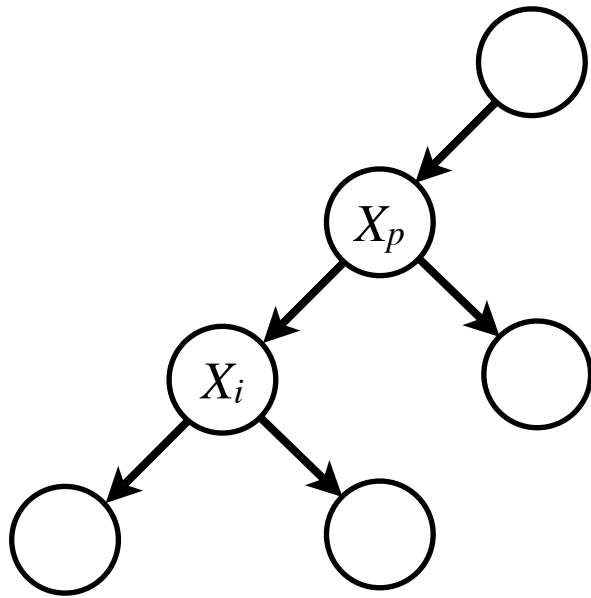
- We now have a way to recursively compute all the $\lambda(X_i)$'s.
- Each node in the network passes a little “ λ message” to its parent.



The other half of the problem

- Remember $P(X_i | E) = \alpha \pi(X_i) \lambda(X_i)$. Now that we have all the $\lambda(X_i)$'s, what about the $\pi(X_i)$'s?
 - recall: $\pi(X_i) = P(X_i | E_i^+)$
- At the root node X_r , E_r^+ is the null set, so $\pi(X_r) = P(X_r)$.
 - since we also know $\lambda(X_i)$, we can compute the final $P(X_r | E)$!
- So for an arbitrary X_i with parent X_p , we can recursively compute $\pi(X_i)$ from $\pi(X_p)$ and/or $P(X_p | E)$.
- How do we do this?

Computing $\pi(X_i)$

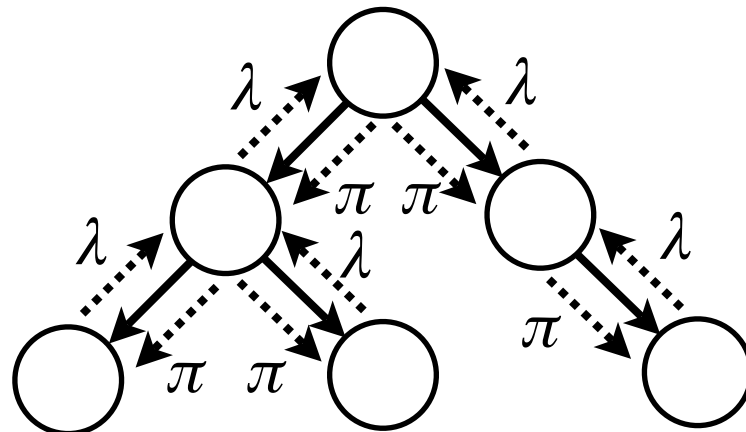


$$\begin{aligned}\pi(X_i) &= P(X_i | E_i^+) \\ &= \sum_j P(X_i, X_p = j | E_i^+) \\ &= \sum_j P(X_i | X_p = j, E_i^+) P(X_p = j | E_i^+) \\ &= \sum_j P(X_i | X_p = j) P(X_p = j | E_i^+) \\ &= \sum_j P(X_i | X_p = j) \frac{P(X_p = j | E)}{\lambda_i(X_p = j)} \\ &= \sum_j P(X_i | X_p = j) \pi_i(X_p = j)\end{aligned}$$

where $\pi_i(X_p)$ is defined as $\frac{P(X_p|E)}{\lambda_i(X_p)}$.

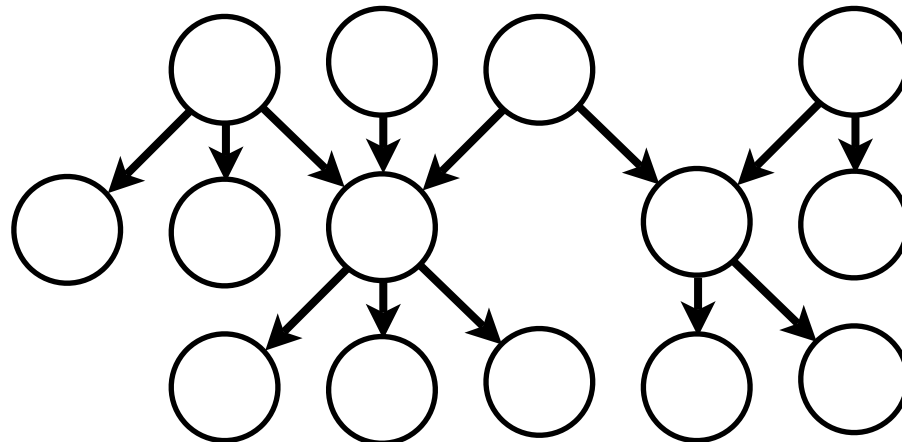
Belief Propagation Wrap-up

- Using a top down strategy, we can compute, for all i , the $P(X_i | E)$ and, in turn, the $\pi(X_i)$ for the children of X_i .
- Nodes can be thought of as autonomous processors passing λ and π messages to their neighbors.
- What if you have a conjunctive distribution, e.g., $P(A, B | C)$ instead of just marginal distributions $P(A | C)$ and $P(B | C)$?
 - Use the chain rule: $P(A, B | C) = P(A | C)P(B | A, C)$
 - Each of these latter probabilities can be computed using belief propagation.



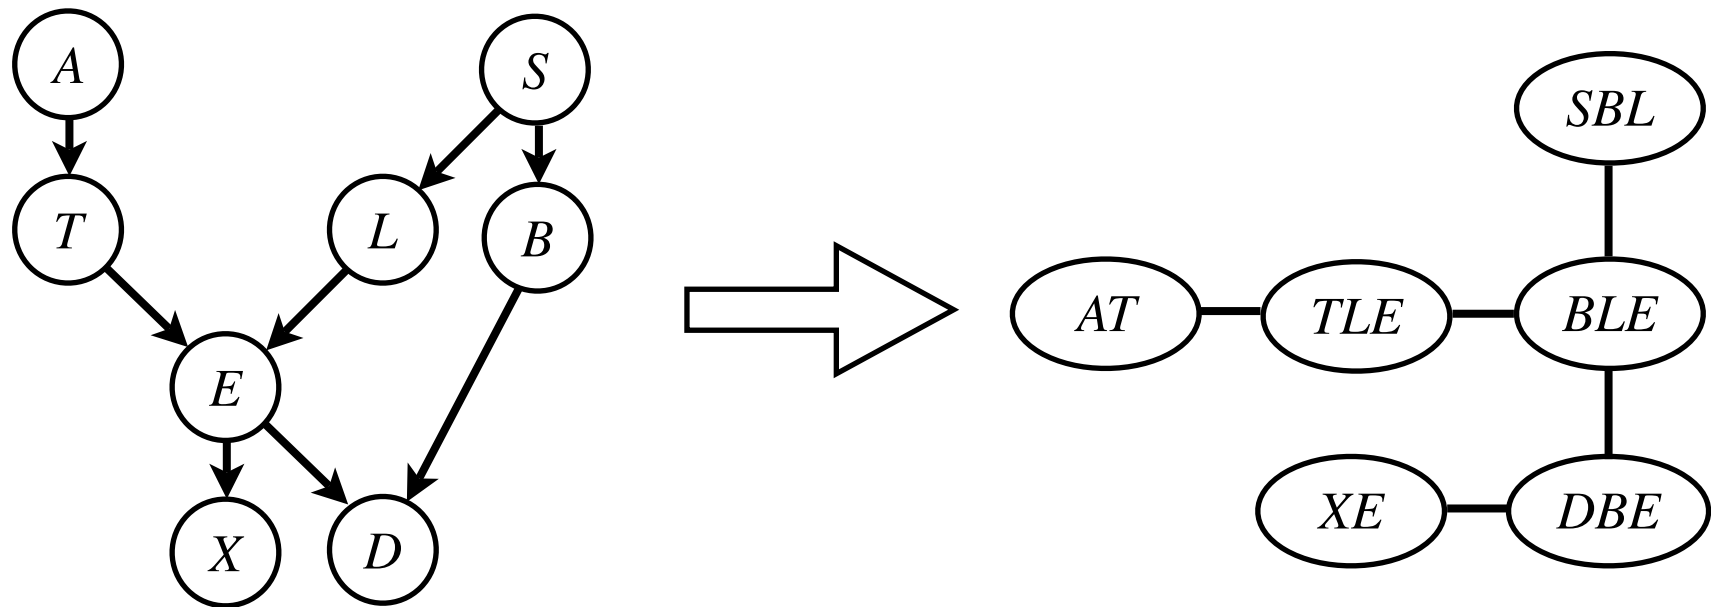
Generalizing Belief Propagation

- Great! For tree structured graphs, we have an inference algorithm with time complexity linear in the number of nodes.
- All this is fine, but what if you don't have a tree structured graph?
- The method we discussed can be generalized to *polytrees*:
 - undirected versions of the graphs are still trees, but nodes can have more than one parent.
 - **do not contain any undirected cycles.**



Dealing with Cycles

- Can deal with undirected cycles in a graph by the *junction tree algorithm*.
- An arbitrary directed acyclic graph can be transformed via some graph theoretic procedure into a tree structure with nodes containing clusters of variables (cliques). Belief propagation can be performed on the new tree.



- Note: In the worst case, the junction tree nodes must take on exponentially many combinations of values, but can work well in practice.

Approximate Inference Methods

- For densely connected graphs, exact inference may be intractable.
- There are 3 widely used approximation schemes
 - Pretend the graph is a tree: “loopy belief propagation”
 - Markov chain Monte Carlo (MCMC): simulate lots of values of the variables and use the sample frequencies to approximate probabilities.
 - Variational inference: Approximate the graph with a simplified the graph (with no cycles) and fit a set of “variational” parameters such that the probabilities computed with the simple structure are “close” to those in the original graph.

Acknowledgments

- The material presented here is taken from tutorials, notes and lecture slides from Yoshua Bengio, Christopher Bishop, Andrew Moore, Tom Mitchell and Scott Davies.

Note to other teachers and users of these slides. Andrew and Scott would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: <http://www.cs.cmu.edu/~awm/tutorials>. Comments and corrections gratefully received.