



C.M.Bishop: Pattern Recognition and Machine Learning. Springer, 2006.

Approximate Energy Minimisation by Belief Propagation

COMPSCI 773 S1 T

VISION GUIDED CONTROL

A/P Georgy Gimel'farb







Graph Models: Joint P.D.

- **Joint probability distribution** (p.d.) over a set of variables: a product of factors (functions depending each on a subset of variables \mathbf{x}_s)

$$\mathbf{x} = (x_1, \dots, x_n) \Rightarrow p(\mathbf{x}) = \prod_s f_s(\mathbf{x}_s)$$

- **Factor graph:** a node for every variable x_i and an additional node for each factor \mathbf{x}_s in the joint p.d.
 - Factor graph for a p.d. $p(\mathbf{x}) = f_a(x_1, x_2)f_b(x_1, x_2)f_c(x_2, x_3)f_d(x_3)$





Factor Graphs



- Factor graphs convey more detailed information about the underlying factorisation of a p.d. than models with only the variable nodes
- Factor graphs are **bipartite**: two distinct kinds of nodes; all links go between nodes of opposite type

– The same undirected or directed graph has many different factor graphs

$$p(\mathbf{x}) = f(x_1, x_2, x_3)$$

$$= f_a(x_1, x_2, x_3) f_b(x_2, x_3)$$

$$p(\mathbf{x}) = f(x_1, x_2, x_3) = \underbrace{f_a(x_1)}_{p_1(x_1)} \underbrace{f_b(x_2)}_{p_2(x_2)} \underbrace{f_c(x_1, x_2, x_3)}_{p(x_3|x_1, x_2)}$$





The Sum-product Algorithm

- Evaluating local marginal probability distributions over a single node or a subset of nodes


- Joint probability distribution: $p(\mathbf{x}) \equiv p(x_1, \dots, x_n) = \prod_s f_s(\mathbf{x}_s)$

- Marginal distribution of a single variable:

$$p(x_i) = \sum_{\mathbf{x} \setminus x_i} p(\mathbf{x}) \equiv \sum_{x_1} \dots \sum_{x_{i-1}} \sum_{x_{i+1}} \dots \sum_{x_n} p(x_1, \dots, x_i, \dots, x_n)$$

- *Summing* - for the assumed discrete variables
- *Integration* - for continuous variables (e.g. linear-Gaussian models)

- **Belief propagation** for exact inference on **digraphs without loops** is a special case of the **sum-product** algorithm





The Sum-product Algorithm

- Original graph - an undirected or directed tree, or a polytree
 - **The corresponding factor graph has a tree structure**
- To compute the marginal $p(x) = \sum_{\mathbf{x} \setminus x} p(\mathbf{x})$:
 - Substitute for $p(\mathbf{x})$ using the factor graph expression $p(\mathbf{x}) = \prod_s f_s(\mathbf{x}_s)$
 - Then interchange summations and products for obtaining a computationally efficient algorithm:

$$p(x) = \sum_{\mathbf{x} \setminus x} \prod_s f_s(\mathbf{x}_s) \Rightarrow \prod_{s \in \text{Nei}(x)} \sum_{\mathbf{x}_s} f_s(x, \mathbf{x}_s)$$







Computing Marginals

- Joint distribution of n variables $p(x_1, x_2, \dots, x_n)$; x_i in $[0, 1, \dots, X-1]$
- Marginal distribution of a variable x_i : generally, X^{n-1} operations!

$$p(x_i) = \sum_{x_1=0}^{X-1} \cdots \sum_{x_{i-1}=0}^{X-1} \sum_{x_{i+1}=0}^{X-1} \cdots \sum_{x_n=0}^{X-1} p(x_1, \dots, x_n)$$

– A very simple particular case:

$$p(x_1, \dots, x_n) = f_1(x_1, x_2) f_2(x_2, x_3) \cdots f_{n-1}(x_{n-1}, x_n)$$






Computing a Marginal...

- $p(x_i)$ for $p(x_1, \dots, x_n) = f_1(x_1, x_2) f_2(x_2, x_3) \cdots f_{n-1}(x_{n-1}, x_n)$:
only $(N-1)X^2$ operations by interchanging the sums and the products:

$$\begin{aligned}
 p(x_i) = & \underbrace{\sum_{x_{i-1}=0}^{X-1} f_{i-1}(x_{i-1}, x_i)}_{\varphi_{i-1}(x_i)} \cdots \underbrace{\sum_{x_2=0}^{X-1} f_2(x_2, x_3)}_{\varphi_2(x_3)} \underbrace{\sum_{x_1=0}^{X-1} f_1(x_1, x_2)}_{\varphi_1(x_2)} \\
 & \times \underbrace{\sum_{x_{i+1}=0}^{X-1} f_i(x_i, x_{i+1}) \cdots \sum_{x_{n-2}=0}^{X-1} f_{n-2}(x_{n-2}, x_{n-1})}_{\psi_{n-2}(x_{n-2})} \underbrace{\sum_{x_n=0}^{X-1} f_{n-1}(x_{n-1}, x_n)}_{\psi_{n-1}(x_{n-1})} \\
 & \underbrace{\hspace{15em}}_{\psi_i(x_i)}
 \end{aligned}$$

$$x_1 \Rightarrow x_i$$

$$x_i \Leftarrow x_n$$

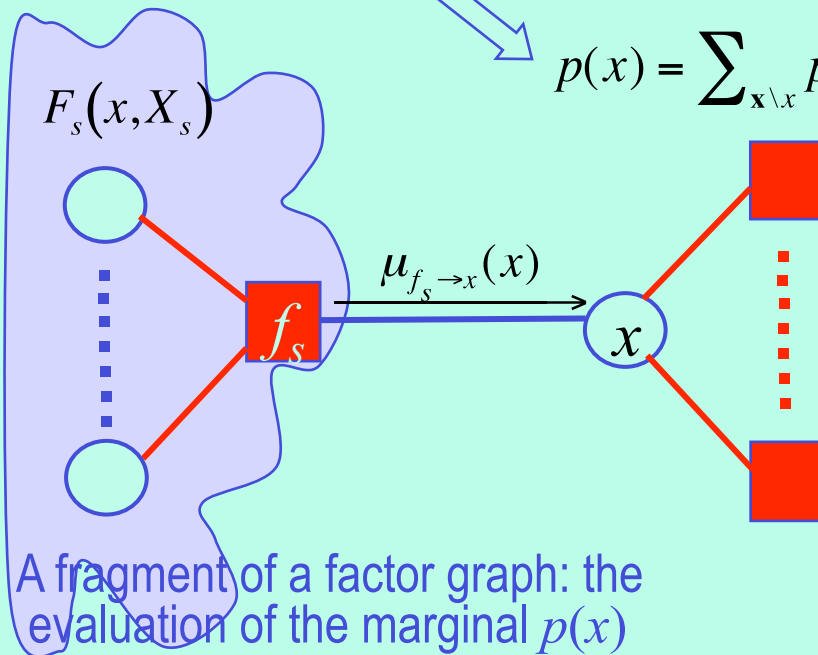




The Sum-product Algorithm


$$p(\mathbf{x}) = \prod_{s \in \text{Nei}(x)} F_s(x, X_s)$$

$$p(x) = \sum_{\mathbf{x} \setminus x} p(\mathbf{x})$$



Messages from f_s to x : $\mu_{f_s \rightarrow x}(x) \equiv \sum_{X_s} F_s(x, X_s)$

- Partitioning the factors in the joint distribution $p(\mathbf{x})$ into groups due to the tree structure of the graph:
 - One group is associated with each of the factor nodes that is a *neighbour* of the variable node x
 - $\text{Nei}(x)$ - the set of factor nodes f_s that are neighbours of x
 - X_s - the set of all variables in a subtree connected to the variable node x via the factor node f_s
 - $F_s(x, X_s)$ - the product of all the factors in the group associated with the factor node f_s

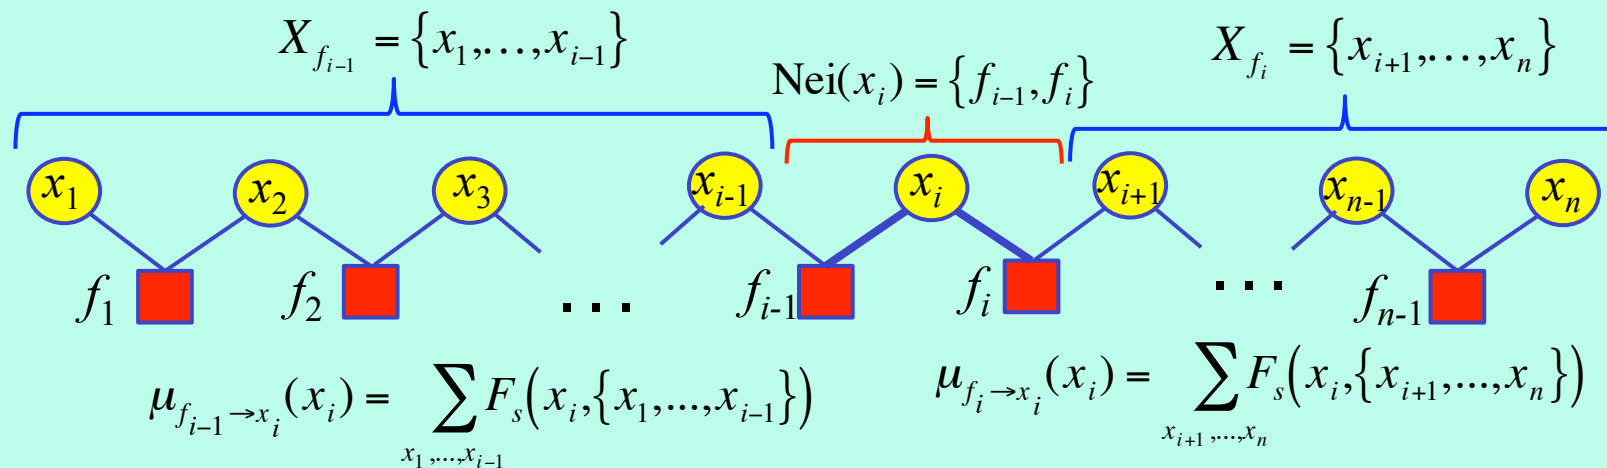




Computing a Marginal...

Particular case:
$$p(\mathbf{x}) = \prod_{s \in \text{Nei}(x_i) = \{f_{i-1}, f_i\}} F_s(x_i, X_s) = \underbrace{\prod_{k=1}^{i-1} f_k(x_k, x_{k+1})}_{F_{f_{i-1}}(x_i, X_{f_{i-1}})} \underbrace{\prod_{k=i}^{n-1} f_k(x_k, x_{k+1})}_{F_{f_i}(x_i, X_{f_i})}$$

$$p(x_i) = \prod_{s \in \text{Nei}(x_i)} \sum_{X_s} F_s(x_i, X_s) = \mu_{f_{i-1} \rightarrow x_i}(x_i) \mu_{f_i \rightarrow x_i}(x_i)$$





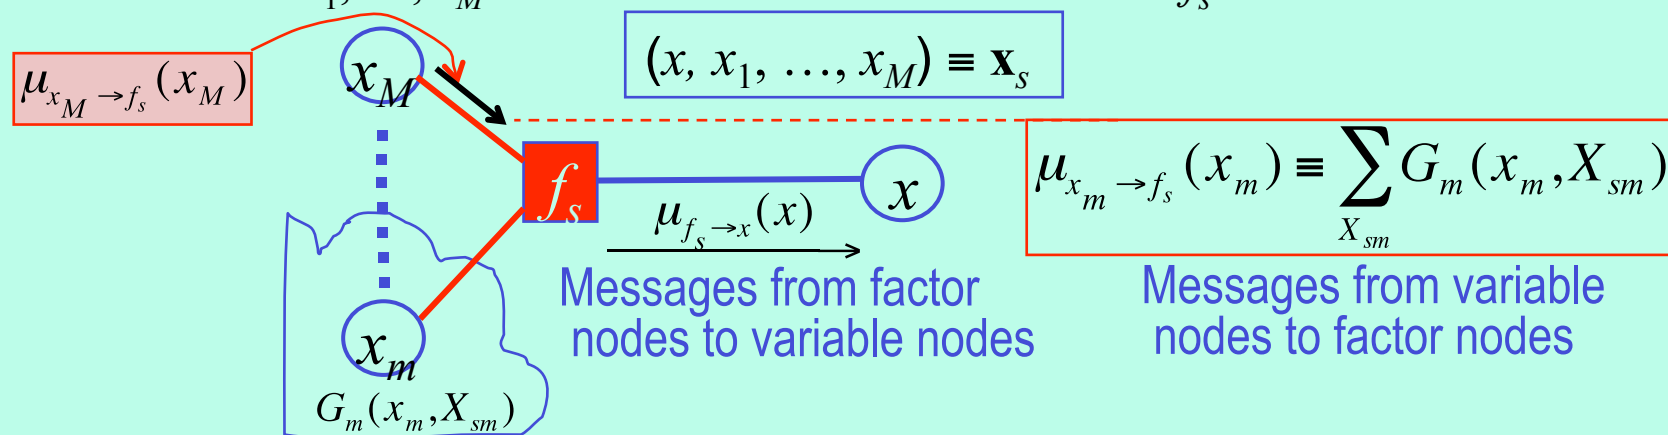
The Sum-product Algorithm

- The goal marginal is the product of all the incoming **messages**:

$$p(x) = \sum_{\mathbf{x} \setminus x} p(\mathbf{x}) = \sum_{\mathbf{x} \setminus x} \prod_{s \in \text{Nei}(x)} F_s(x, X_s) = \prod_{s \in \text{Nei}(x)} \left[\sum_{X_s} F_s(x, X_s) \right] = \prod_{s \in \text{Nei}(x)} \mu_{f_s \rightarrow x}(x)$$

- Each factor $F_s(x, X_s)$ is described by a factor (sub)-graph and so can itself be factorised: $F_s(x, X_s) = f_s(x, x_1, \dots, x_M) G_1(x_1, X_{s1}) \dots G_1(x_M, X_{sM})$

- x_1, \dots, x_M are the variables associated with factor f_s in addition to x





$$\mu_{x_m \rightarrow f_s}(x_m) \equiv \sum_{X_{sm}} G_m(x_m, X_{sm})$$

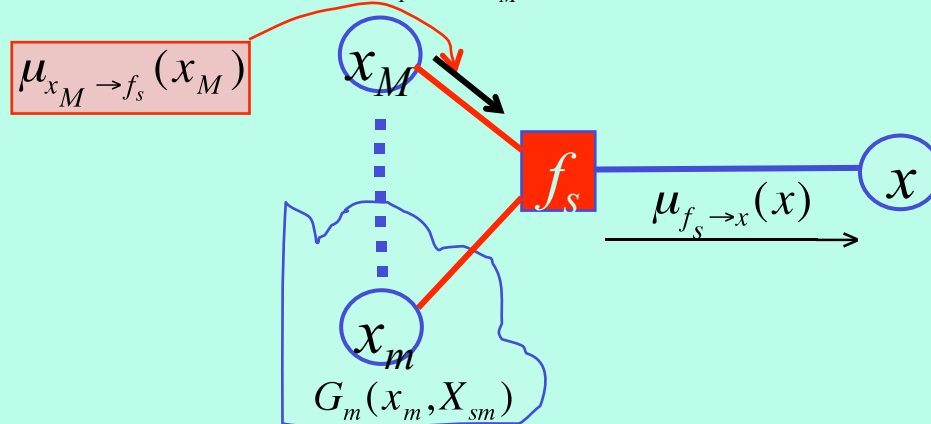
The Sum-product Algorithm

- Computing the messages from factor nodes to variable nodes

[$\text{Nei}(f_s)$ - the set of variable nodes being neighbours of the factor node f_s


$\text{Nei}(f_s) \setminus x$ - the same set but without node x]:

$$\begin{aligned} \mu_{f_s \rightarrow x}(x) &\equiv \sum_{X_s} F_s(x, X_s) = \sum_{x_1} \cdots \sum_{x_M} f_s(x, x_1, \dots, x_M) \prod_{m \in \text{Nei}(f_s) \setminus x} \left[\sum_{X_{sm}} G_m(x_m, X_{sm}) \right] \\ &= \sum_{x_1} \cdots \sum_{x_M} f_s(x, x_1, \dots, x_M) \prod_{m \in \text{Nei}(f_s) \setminus x} \mu_{x_m \rightarrow f_s}(x_m) \end{aligned}$$



Evaluating the message sent by a factor node to a variable node along their connecting link:

- Take the product of the incoming messages along all other links coming to the factor node,
- Multiply by the factor associated with that node, and
- Marginalise over all the variables associated with the incoming messages






$$\mu_{x_m \rightarrow f_s}(x_m) \equiv \sum_{X_{sm}} G_m(x_m, X_{sm})$$

The Sum-product Algorithm

- A factor node can send a message to a variable node once it has received incoming messages from all other neighbouring variable nodes




- Evaluating the messages from variable nodes to factor nodes - again by the (sub)-graph factorization

- The term $G_m(x_m, X_{sm})$ associated with node x_m is given by a product of terms $F_l(x_m, X_{ml})$ each associated with one of the factor nodes f_l that is linked to node x_m (excluding node f_s):

$$G_m(x_m, X_{sm}) = \prod_{l \in \text{Nei}(x_m) \setminus f_s} F_l(x_m, X_{ml})$$

- The product: over all neighbours of node x_m except for node f_s
- Each of the factors $F_l(x_m, X_{ml})$: a subtree of the original graph of the same kind as the joint $p(\mathbf{x})$





$$\mu_{f_s \rightarrow x}(x) \equiv \sum_{X_s} F_s(x, X_s)$$

$$\mu_{x_m \rightarrow f_s}(x_m) \equiv \sum_{X_{sm}} G_m(x_m, X_{sm})$$

$$G_m(x_m, X_{sm}) = \prod_{l \in \text{Nei}(x_m) \setminus f_s} F_l(x_m, X_{ml})$$


The Sum-product Algorithm

- Evaluating the message sent by a variable node to an adjacent factor node along the connecting link:

- The product of the incoming messages along all of the other links:

$$\mu_{x_m \rightarrow f_s}(x_m) = \prod_{l \in \text{Nei}(x_m) \setminus f_s} \left[\sum_{X_{ml}} F_l(x_m, X_{ml}) \right] = \prod_{l \in \text{Nei}(x_m) \setminus f_s} \mu_{f_l \rightarrow x_m}(x_m)$$

- Any variable node having only two neighbours simply passes messages through unchanged
- A variable node can send a message to a factor node once it has received incoming messages from all other neighbouring factor nodes
- The marginal for variable node x is given by the product of incoming messages along all of the links arriving at that node
 - Each of these messages can be computed recursively in terms of other messages by viewing x as the root of the tree and starting at the leaf nodes






$$p(x) = \prod_{s \in \text{Nei}(x)} \mu_{f_s \rightarrow x}(x)$$

$$p(\mathbf{x}_s) = f_s(\mathbf{x}_s) \prod_{i \in \text{Nei}(f_s)} \mu_{x_i \rightarrow f_s}(x_i)$$

The Sum-product Algorithm


- If a leaf node is a variable node, then $\mu_{x \rightarrow f}(x) = 1$
- If the leaf node is a factor node, then $\mu_{f \rightarrow x}(x) = f(x)$
- General algorithm:
 - Arbitrarily pick any (variable or factor) node and designate it as the root
 - Propagate messages from the leaves to the root until the root node will have received messages from all of its neighbours
 - Once this message propagation is complete, then propagate messages from the root to all of its neighbours and further along all of the links outwards from the root all the way to the leaves
 - Now, a message will pass in both directions across every link in the graph, and every node will receive a message from all of its neighbours
 - The marginal distribution is readily calculated for every variable in the graph because every variable node has received messages from all its neighbours





The Sum-product Algorithm

- After one message has passed in each direction across each link, the marginals are evaluated as $p(x) = \prod_{s \in \text{Nei}(x)} \mu_{f_s \rightarrow x}(x)$
- Normalisation of the distribution
 - If the factor graph is derived from a directed graph, then the joint distribution is already correctly normalised
 - So the computed marginals will be similarly normalised
 - For an undirected graph, in general there is an unknown factor $1/Z$
 - The sum-product algorithm is to be run first with an unnormalised version of the joint distribution in order to find the unnormalised marginals
 - The factor $1/Z$ is then easily obtained by normalising any one of these marginals (a computationally efficient approach because the normalisation is done over a single variable rather than over the entire set of variables)

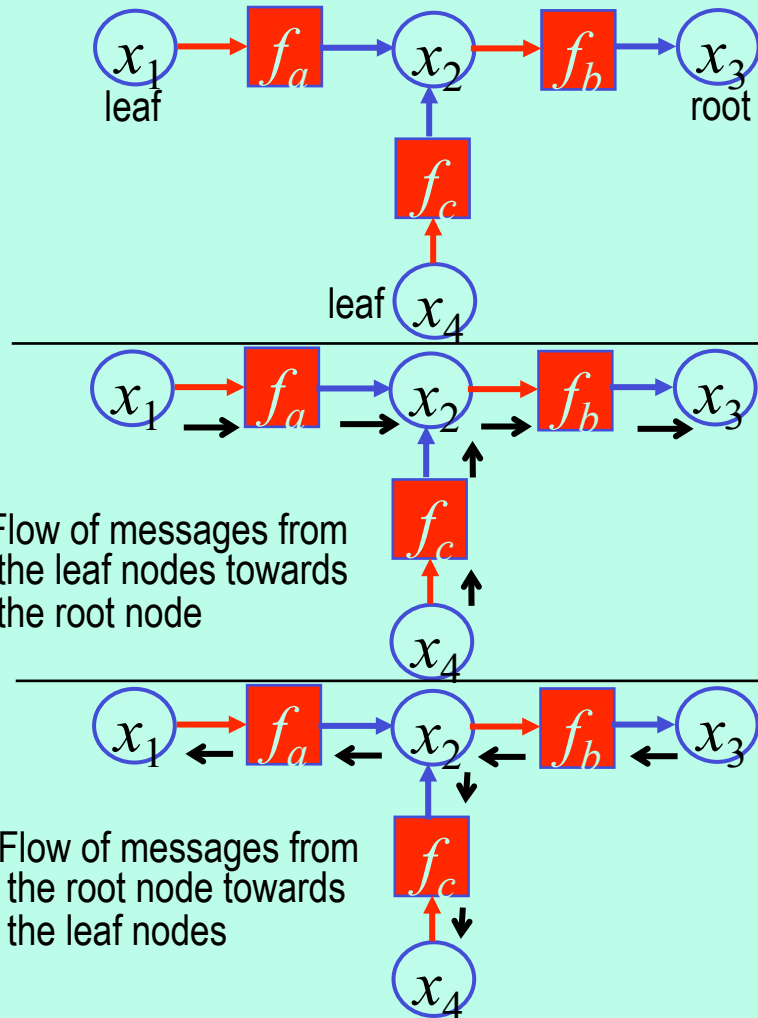




$$p(x) = \prod_{s \in \text{Nei}(x)} \mu_{f_s \rightarrow x}(x)$$

$$p(\mathbf{x}_s) = f_s(\mathbf{x}_s) \prod_{i \in \text{Nei}(f_s)} \mu_{x_i \rightarrow f_s}(x_i)$$

The Sum-product Algorithm: An Example



- Unnormalised joint distribution:

$$\tilde{p}(\mathbf{x}) = f_a(x_1, x_2) f_b(x_2, x_3) f_c(x_2, x_4)$$

Starting with the leaf nodes:

$$\mu_{x_1 \rightarrow f_a}(x_1) = 1; \quad \mu_{f_a \rightarrow x_2}(x_2) = \sum_{x_1} f_a(x_1, x_2)$$

$$\mu_{x_4 \rightarrow f_c}(x_4) = 1; \quad \mu_{f_c \rightarrow x_2}(x_2) = \sum_{x_4} f_c(x_2, x_4)$$

$$\mu_{x_2 \rightarrow f_b}(x_2) = \mu_{f_a \rightarrow x_2}(x_2) \mu_{f_c \rightarrow x_2}(x_2)$$

$$\mu_{f_b \rightarrow x_3}(x_3) = \sum_{x_2} f_b(x_2, x_3) \mu_{x_2 \rightarrow f_b}(x_2)$$

Propagating messages from the root node out to the leaf nodes


$$\mu_{x_3 \rightarrow f_b}(x_1) = 1; \quad \mu_{f_b \rightarrow x_2}(x_2) = \sum_{x_3} f_b(x_2, x_3)$$

$$\mu_{x_2 \rightarrow f_a}(x_2) = \mu_{f_b \rightarrow x_2}(x_2) \mu_{f_c \rightarrow x_2}(x_2);$$

$$\mu_{f_a \rightarrow x_1}(x_1) = \sum_{x_2} f_a(x_1, x_2) \mu_{x_2 \rightarrow f_a}(x_2)$$

$$\mu_{x_2 \rightarrow f_c}(x_2) = \mu_{f_a \rightarrow x_2}(x_2) \mu_{f_b \rightarrow x_2}(x_2)$$

$$\mu_{f_c \rightarrow x_4}(x_4) = \sum_{x_2} f_c(x_2, x_4) \mu_{x_2 \rightarrow f_c}(x_2)$$





The Max-sum Algorithm

- An application of *dynamic programming* in the graphical models
- Find the set of values that jointly have the max probability:

$$\mathbf{x}^{\max} = \arg \max_{\mathbf{x}} p(\mathbf{x}) \Leftrightarrow p(\mathbf{x}^{\max}) = \max_{\mathbf{x}} p(\mathbf{x})$$


- An example for a chain of nodes:

$$p(\mathbf{x}) = \frac{1}{Z} \psi_{1,2}(x_1, x_2) \cdot \dots \cdot \psi_{N-1,N}(x_{N-1}, x_N)$$

$$\max_{\mathbf{x}} p(\mathbf{x}) = \frac{1}{Z} \max_{x_1} \left[\max_{x_2} \psi_{1,2}(x_1, x_2) \left[\dots \max_{x_N} \psi_{N-1,N}(x_{N-1}, x_N) \right] \right]$$

$$\ln \left(\max_{\mathbf{x}} p(\mathbf{x}) \right) = \max_{\mathbf{x}} \ln p(\mathbf{x}) \equiv \max_{\mathbf{x}} \left[\ln \psi_{1,2}(x_1, x_2) + \dots + \ln \psi_{N-1,N}(x_{N-1}, x_N) \right]$$

$$\max_{\mathbf{x}} \ln p(\mathbf{x}) = -\ln Z + \max_{x_1} \left[\max_{x_2} \ln \psi_{1,2}(x_1, x_2) + \left[\dots + \max_{x_N} \ln \psi_{N-1,N}(x_{N-1}, x_N) \right] \right]$$





The Max-sum Algorithm


- Similarly to the sum-product algorithm, the max-sum algorithm is readily written down in terms of message passing by replacing ‘sum’ with ‘max’ and products with sums of logarithms:

$$\mu_{f \rightarrow x}(x) = \max_{x_1, \dots, x_M} \left[\ln f(x, x_1, \dots, x_M) + \sum_{m \in \text{ne}(f) \setminus x} \mu_{x_m \rightarrow f}(x_m) \right]$$

$$\mu_{x \rightarrow f}(x) = \sum_{l \in \text{ne}(x) \setminus f} \mu_{f_l \rightarrow x}(x)$$

- The initial messages sent by the leaf nodes:

$$\begin{cases} \mu_{f \rightarrow x}(x) = 0 \\ \mu_{x \rightarrow f}(x) = \ln f(x) \end{cases}$$






The Max-sum Algorithm

- The maximum probability at the root node:

$$p^{\max} = \max_x \left[\sum_{s \in \text{ne}(x)} \mu_{f_s \rightarrow x}(x) \right]; \quad x^{\max} = \arg \max_x \left[\sum_{s \in \text{ne}(x)} \mu_{f_s \rightarrow x}(x) \right]$$



- The maximum of the joint probability is found by propagating messages from leaves to an arbitrary chosen root node
- The result is the same irrespective of which node is chosen as the root
- But the propagation of the messages back from the root to the leaves (with maximisation instead of summing) fails to return \mathbf{x}^{\max}
 - Individual variable values found by maximising the product of messages at each node may belong to different maximizing configurations \mathbf{x}^{\max}
 - An overall configuration may no longer correspond to a maximum





The Max-sum Algorithm

- To resolve the problem: a different kind of message passing from the root node to the leaves
 - Example:** a chain of N variables x_1, \dots, x_N each having K states





The Max-sum Algorithm

- Suppose node x_N is taken to be the root node

Phase 1: Propagate messages from the leaf node x_1 to the root node:

$$\mu_{x_n \rightarrow f_{n,n+1}}(x_n) = \mu_{f_{n-1,n} \rightarrow x_n}(x_n); \quad \mu_{x_1 \rightarrow f_{1,2}}(x_1) = 0$$


$$\mu_{f_{n-1,n} \rightarrow x_n}(x_n) = \max_{x_{n-1}} \left[\ln f_{n-1,n}(x_{n-1}, x_n) + \mu_{x_{n-1} \rightarrow f_{n-1,n}}(x_{n-1}) \right]$$

$$\phi(x_n) = \arg \max_{x_{n-1}} \left[\ln f_{n-1,n}(x_{n-1}, x_n) + \mu_{x_{n-1} \rightarrow f_{n-1,n}}(x_{n-1}) \right]$$

- There could be several values x_{n-1} giving above the same maximum

Phase 2: Backtracking - propagate a message back down the chain:

$$x_N^{\max} = \arg \max_{x_N} [\mu_{f_{N-1,N} \rightarrow x_N}(x_N)]; \quad x_{n-1}^{\max} = \phi(x_n^{\max})$$





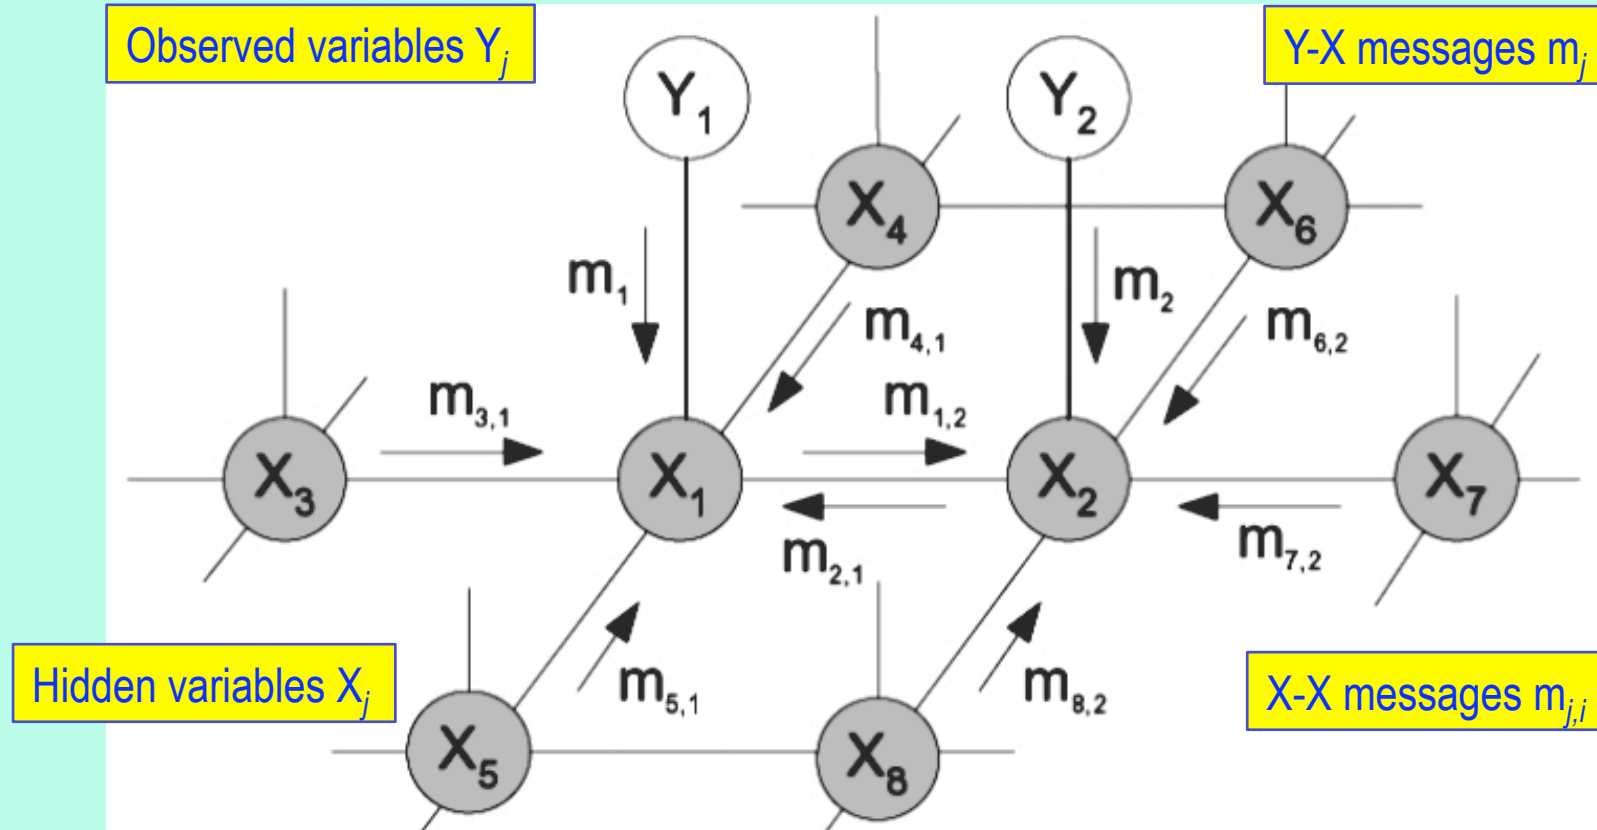
The Max-sum Algorithm

- To have a global maximum of a joint probability distribution:
 1. Keep track of the maximising states during the forward pass ($\phi(x_n)$)
 2. Use back-tracking to find a consistent solution
- Extension to a general tree-structured factor graph:
 - If a message is sent from a factor node f to a variable node x , a maximisation is performed over all other variable nodes x_1, \dots, x_M that are neighbours of that factor node
 - When this maximisation is performed, we keep a record of which values of the variables x_1, \dots, x_M gave rise to the maximum
 - Then in the backtracking step, having found x^{\max} , we use the stored values to assign consistent maximising states $x_1^{\max}, \dots, x_M^{\max}$
- Important application: the *Viterbi* algorithm (to find a HMM model)






MRF Model for Stereo



Posterior model:
$$P(\mathbf{X} | \mathbf{Y}) \propto \prod_j \left(\psi_j(\mathbf{X}_j, Y_j) \prod_{i \in \text{Neib}(j)} \varphi_{ji}(\mathbf{X}_j, \mathbf{X}_i) \right)$$





Max-sum Algorithm for Stereo

- Log-likelihood: $\log P(\mathbf{X} | \mathbf{Y}) \propto \sum_j \left(\log \psi_j(\mathbf{X}_j, \mathbf{Y}_j) + \sum_{i \in \text{Neib}(j)} \log \varphi_{ji}(\mathbf{X}_j, \mathbf{X}_i) \right)$
 - Simplification: $m_{ji}(\mathbf{X}_j, \mathbf{X}_i) = m_{ji}(\mathbf{X}_i)$; $m_j(\mathbf{X}_j, \mathbf{Y}_j) = m_j(\mathbf{X}_j)$


- Max-sum BP:

1. Initialise all messages $m_{ji}(\mathbf{X}_i) = 0$ and $m_j(\mathbf{X}_j) = \log \psi_j(\mathbf{X}_j, \mathbf{Y}_j)$
2. Update messages $m_{ji}(\mathbf{X}_i)$ iteratively for $j = 1, \dots, J$

$$m_{ji}^{[t+1]}(\mathbf{X}_i) \leftarrow \max_{x_j} \left\{ \log \varphi_{ji}(\mathbf{X}_j, \mathbf{X}_i) + m_j^{[t]}(\mathbf{X}_j) + \sum_{k \in \text{Neib}(j) \setminus i} m_{kj}^{[t]}(\mathbf{X}_j) \right\}$$

3. Compute beliefs:


$$b_j(\mathbf{X}_j) \leftarrow m_j^{[t]}(\mathbf{X}_j) + \sum_{k \in \text{Neib}(j)} m_{kj}^{[t]}(\mathbf{X}_j); \quad \mathbf{X}_j^{\text{MAP}} = \arg \max_{\mathbf{X}} b_j(\mathbf{X})$$





Exact Inference in General Graphs

- The MAP solution $\mathbf{y}^* = \arg \max_{\mathbf{y}} p(\mathbf{y}|\mathbf{x})$ can be found by special random sampling from $p(\mathbf{y}|\mathbf{x})$ called *simulated annealing*
 - But its convergence is too slow in most practical cases
- The *sum-product* and *max-sum* algorithms provide efficient and exact solutions to inference problems in tree-structured graphs
- For many practical applications, graphs have loops
 - The message passing framework can be generalised to arbitrary graphs topologies to give an exact inference procedure
 - But in the case of discrete variables, its computational complexity grows exponentially with the maximum number of interdependent variables
 - Thus for many problems of practical interest, it is not feasible to use exact inference, so that effective approximation methods have to be exploited





Approximate Inference

- ***Loopy belief propagation*** (LBP)
 - Is possible because the message passing rules for the sum-product algorithm are purely local
 - But because now the graph has cycles, information can flow many times around the graph
 - For some models, the algorithm converges, whereas for others it will not
 - **Message passing schedule**: when each node should send a message across link from the node after receiving messages from all other links
 - Transmit only *pending* messages: after a node receives a message on one of its links (convergence to the exact marginal if no more pending messages)
- ***Graph-cut-based*** algorithms
 - Exact inference only for binary vector signals \mathbf{x} and decisions \mathbf{y}

