

Solution to Problem Set 4

Fall 2012

Issued: Tues. Oct. 2, 2012 **Due:** Tues. Oct. 16, 2012

Reading: Chapters 4, 17

Problem 4.1

Sum-product algorithm: Consider the sum-product algorithm on an undirected tree with compatibility functions ψ_s and ψ_{st} . For the tree in Figure 1(a), suppose that we have ternary random variables (i.e., $X_s \in \{-1, 0, 1\}$), where the compatibility functions are of the form:

$$\psi_{st}(x_s, x_t) = \begin{bmatrix} a & b & b \\ b & a & b \\ b & b & a \end{bmatrix}, \quad \psi_s(x_s) = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \text{ for } s \text{ odd and, } \psi_s(x_s) = \begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix} \text{ for } s \text{ even}$$

(Here entry (i, j) in the 3×3 matrix notation for ψ_{st} gives the value $\psi_{st}(i, j)$, whereas entry i in the 3-vector ψ_s gives the number $\psi_s(i)$).

Throughout this problem, use the serial ordering of message-passing, in which node s updates its message to t only when it has received all other incoming messages.

- (a) Implement the sum-product algorithm for a general tree. Please document (meaning describe in comments within the code) what each step is doing, and hand in your documented code. Use it to compute marginal distributions for the tree in Figure 1(a), using the specified compatibility functions with
- (i) $a = 1$ and $b = 0.5$
 - (ii) $a = 1$ and $b = 2$.

Solution: The marginals for $a = 1, b = 0.5$ are:

	$\mathbb{P}(X_s) = -1$	$\mathbb{P}(X_s) = 0$	$\mathbb{P}(X_s) = 1$
X_1	0.1617	0.2712	0.5671
X_2	0.4220	0.1435	0.4345
X_3	0.1639	0.2767	0.5594
X_4	0.5058	0.1422	0.3520
X_5	0.1868	0.2889	0.5243
X_6	0.4376	0.1672	0.3952

And the marginals for $a = 1$, $b = 2$ are:

	$\mathbb{P}(X_s) = -1$	$\mathbb{P}(X_s) = 0$	$\mathbb{P}(X_s) = 1$
X_1	0.1690	0.3630	0.4680
X_2	0.5395	0.1835	0.2769
X_3	0.1704	0.3695	0.4601
X_4	0.4461	0.1920	0.3619
X_5	0.1473	0.3560	0.4967
X_6	0.5344	0.1635	0.3021

The code implementing the sum product algorithm is attached below, along with the script to run the algorithm on the tree given in this problem.

```

1 function marginals = SumProduct(adj, psi_s, psi_st)
2 % This function implements the recursive SumProduct algorithm
3 % (without evidence) as outlined in the course reader,
4 % Figure 4.5.
5 %
6 % Input:
7 % * adj      = n-by-n matrix with entries in {0, 1}, representing
8 %             the adjacency matrix of the tree.
9 % * psi_s    = d-by-n matrix whose columns represent the vertex
10 %            potentials. We assume each random variable X_s can
11 %            take d possible values.
12 % * psi_st   = n-by-n cell, whose (s,t)-entry is a d-by-d matrix
13 %            representing the potential function for edge (s,t)
14 %            in the tree. For pairs (s,t) that are not edges
15 %            in the tree, psi_st(s,t) is empty.
16 %
17 % Output:
18 % * marginals = d-by-n matrix whose columns are the marginals
19 %             of the nodes in the tree.
20 %
21
22 n = size(adj,1);    % number of nodes in the graph
23 d = size(psi_s,1); % number of states in each variable
24
25 % Create a list of neighbors for each node
26 neighbors = cell(n,1);
27 for i = 1:n
28     neighbors{i} = find(adj(i,:) == 1);
29 end
30
31 % Initialize messages to be all 1.

```

```

32 % For each edge (i,j), M{i,j} contains the message from node i
33 % to node j, and M{j,i} contains the message from node j to
34 % node i. If (i,j) is not an edge, M{i,j} is empty.
35 M = cell(n,n);
36 for i = 1:n
37     for j = i+1:n
38         if (adj(i,j) == 1)
39             M{i,j} = ones(d,1);
40             M{j,i} = ones(d,1);
41         end
42     end
43 end
44
45 % Arbitrarily choose root
46 root = 1;
47
48 % Collect messages from root's neighbors
49 for e = neighbors{root}
50     Collect(root, e);
51 end
52
53 % Distribute messages to root's neighbors
54 for e = neighbors{root}
55     Distribute(root, e);
56 end
57
58 % Compute marginals for each node
59 marginals = zeros(d,n);
60 for i = 1:n
61     % Multiply the potential of node i with the messages
62     % from its neighbors
63     psi_i = psi_s(:,i);
64     for j = neighbors{i}
65         psi_i = psi_i .* M{j,i};
66     end
67
68     % Then normalize to get the marginal
69     marginals(:,i) = psi_i / sum(psi_i);
70 end
71
72
73 % Helper function for collecting message from node j to
74 % node i. We first collect messages from node j's
75 % neighbors, excluding i, then send message from j to i.
76 function Collect(i,j)
77     for k = setdiff(neighbors{j}, i);
78         Collect(j,k);
79     end
80     SendMessage(j,i);

```

```

81 end
82
83
84 % Helper function for distributing messages from node i
85 % to node j. We first send message from i to j, then
86 % forward the message to node j's neighbors, excluding i.
87 function Distribute(i,j)
88     SendMessage(i,j);
89     for k = setdiff(neighbors{j}, i);
90         Distribute(j,k);
91     end
92 end
93
94
95 % Helper function for sending message from node j to
96 % node i. We first compute the product of psi-j(x-j)
97 % and m_kj(x-j) over all neighbors k of j, excluding i,
98 % then multiply by \psi_{ij}(x_i, x_j) and sum over x-j.
99 function SendMessage(j,i)
100     psi_j = psi_s(:,j);
101     for k = setdiff(neighbors{j},i)
102         psi_j = psi_j .* M{k,j};
103     end
104     M{j,i} = sum(psi_st{i,j} .* repmat(psi_j', [d 1]), 2);
105 end
106
107 end

```

```

1 function marginals = RunSumProduct(a,b)
2 % Run the sum product algorithm for the tree given in Problem 4.1
3 % with the specified vertex and edge potentials.
4
5 % Adjacency matrix of the tree
6 adj = [0 1 1 0 0 0; ...
7         1 0 0 1 1 0; ...
8         1 0 0 0 0 1; ...
9         0 1 0 0 0 0; ...
10        0 1 0 0 0 0; ...
11        0 0 1 0 0 0];
12
13 % Vertex potentials
14 psi_s = [1 3 1 3 1 3; ...
15          2 1 2 1 2 1; ...
16          3 2 3 2 3 2];
17
18 % Edge potentials
19 psi_template = [a b b; b a b; b b a];

```

```

20 psi_st = cell(size(adj));
21 for i = 1:size(adj,1)
22     for j = i+1:size(adj,1)
23         psi_st{i,j} = psi_template;
24         psi_st{j,i} = psi_template;
25     end
26 end
27
28 % Run sum product
29 marginals = SumProduct(adj, psi_s, psi_st);
30
31 % Display results
32 fprintf('Marginals for a = %g, b = %g:\n', a, b);
33 for i = 1:size(marginals,2)
34     fprintf('X%d:  %.4f  %.4f  %.4f\n', i, marginals(:,i));
35 end

```

- (b) For any tree, prove that sum-product correctly computes the singleton marginals in at most tree diameter iterations. (The tree diameter is the length of the longest path between any two nodes.)

Solution:

Convergence: We first prove that the messages stabilize after D iterations, where D is the diameter of the tree. We use induction on the tree diameter D . For $D = 1$, the result is immediate. Consider a graph of diameter D . Let L be the set of all leaf nodes. For each $v \in L$, let w_v be the only element of N_v , which is a singleton set because v is a leaf node. Thus, at the first time step $t = 1$, the (fixed) message v sends to w_v is

$$M_{v \rightarrow w_v}^*(x_{w_v}) = \sum_{x_v} \psi_v(x_v) \psi_{vw_v}(x_v, x_{w_v}).$$

At each subsequent time step, the message each leaf sends to its neighbor is constant, since it does not depend on messages from any other nodes. We construct a new undirected graph G' by stripping each of the leaf nodes from the original graph and redefining the compatibility function for each w_v as

$$\psi'_{w_v}(x_{w_v}) = \psi_{w_v}(x_{w_v}) M_{v \rightarrow w_v}(x_{w_v}).$$

This new graphical model has diameter at most $D - 2$, which follows because the longest path in a tree is between two leaf nodes. Therefore, we can apply the inductive hypothesis to G' . Thus, after at most $D - 2$

time steps, the messages will all converge. Finally, note that replacing the nodes $v \in L$ in G' will not affect any of the fixed-point messages, since the nodes in L always send the same messages to their neighbors. Adding on the first iteration where all leaf nodes send in messages and the last iteration from nodes w_v to the leaf nodes v , we conclude that after D iterations, all messages will converge to a fixed point.

Correctness: We then prove that the fixed-point messages can be used to compute the marginals via

$$p(x_s) \propto \psi_s(x_s) \prod_{t \in N(s)} M_{t \rightarrow s}^*(x_s).$$

We use induction on the number of nodes in the tree. When the tree has one node, the conclusion is trivial. Now let $m \geq 2$, and suppose the conclusion holds for any tree with $m - 1$ nodes. We will show that the result holds for a tree on m nodes, as well.

Renumber the nodes of the tree so that node m is a leaf, with node 1 as its neighbor. In the first iteration, node m sends the message $M_{m \rightarrow 1}^*(x_1) = \sum_{x_m} \psi_m(x_m) \psi_{1m}(x_1, x_m)$ to node 1. Note that by marginalizing out x_m and using the factorization over potentials, we have

$$\begin{aligned} p(x_1, \dots, x_{m-1}) &= \sum_{x_m} p(x_1, \dots, x_m) \\ &\propto M_{m \rightarrow 1}^*(x_1) \prod_{i=1}^{m-1} \psi_i(x_i) \prod_{(j,k) \in E \setminus \{(1,m)\}} \psi_{jk}(x_j, x_k). \end{aligned}$$

In particular, all subsequent messages sent throughout the rest of the graph behave as the sum-product algorithm on the nodes $\{1, \dots, m - 1\}$, with the potential on node 1 replaced by $\psi'_1(x_1) = M_{m \rightarrow 1}^*(x_1) \psi_1(x_1)$. By the induction hypothesis, the messages on all edges between nodes $\{1, \dots, m - 1\}$ converge to the proper messages $M_{t \rightarrow s}^*(x_s)$, such that the marginal probability of each node is proportional to the product of the node potential and all incoming messages.

It remains to show that the sum-product algorithm computes the correct marginal on node m ; i.e., $p(x_m) \propto \psi_m(x_m) M_{1 \rightarrow m}^*(x_m)$. Since $p(x_1)$ is proportional to the product of $\psi_1(x_1)$ and all incoming messages, and $M_{1 \rightarrow m}(x_m)$ is a sum of the product of $\psi_1(x_1) \psi_{1m}(x_1, x_m)$ and the incoming messages from nodes in $N(1) \setminus \{m\}$, we have

$$M_{1 \rightarrow m}^*(x_m) \propto \sum_{x_1} \frac{p(x_1)}{M_{m \rightarrow 1}^*(x_1)} \psi_{1m}(x_1, x_m).$$

Hence, we see that

$$\psi_m(x_m)M_{1 \rightarrow m}^*(x_m) \propto \psi_m(x_m) \sum_{x_1} \frac{p(x_1)}{\sum_{x'_m} \psi_m(x'_m)\psi_{1m}(x_1, x'_m)} \psi_{1m}(x_1, x_m),$$

where we have substituted in the form of $M_{m \rightarrow 1}^*(x_1)$. Finally, note that marginalizing out the factorized form of $p(x)$ with potentials (summing over all variables except x_1 and x_m in the numerator, and all variables except x_1 in the denominator), we have

$$p(x_m|x_1) = \frac{p(x_1, x_m)}{p(x_1)} = \frac{\psi_m(x_m)\psi_{1m}(x_1, x_m)}{\sum_{x'_m} \psi_m(x'_m)\psi_{1m}(x_1, x'_m)}.$$

It follows that

$$\psi_m(x_m)M_{1 \rightarrow m}^*(x_m) \propto \sum_{x_1} p(x_1)p(x_m|x_1) = p(x_m),$$

as wanted.

- (c) Show that for each edge $(s, t) \in E$, the message fixed point M^* can be used to compute the pairwise joint distribution over (x_s, x_t) as follows:

$$p(x_s, x_t) \propto \psi_s(x_s)\psi_t(x_t)\psi_{st}(x_s, x_t) \prod_{u \in N(s) \setminus t} M_{us}^*(x_s) \prod_{u \in N(t) \setminus s} M_{ut}^*(x_t).$$

Solution: Contract the edge (s, t) into a supernode st with vertex potential $\tilde{\psi}_{st}(x_s, x_t) = \psi_s(x_s)\psi_t(x_t)\psi_{st}(x_s, x_t)$, and consider running the sum product algorithm on this new tree. It is easy to see that the fixed-point messages $\tilde{M}_{uv}^*(x_v)$ in this new tree are exactly the same as the old fixed-point messages $M_{uv}^*(x_v)$ if $v \notin \{s, t\}$. Furthermore, if u is a neighbor of st in the new tree, then $\tilde{M}_{u(st)}^*(x_s, x_t) = M_{us}^*(x_s)$ if (s, u) is an edge in the original tree, and $\tilde{M}_{u(st)}^*(x_s, x_t) = M_{ut}^*(x_t)$ otherwise. Therefore, using the result of part (b) for the marginal of the supernode st , we conclude that

$$\begin{aligned} p(x_s, x_t) &\propto \tilde{\psi}_{st}(x_s, x_t) \prod_{u \in N(st)} \tilde{M}_{u(st)}^*(x_s, x_t) \\ &= \psi_s(x_s)\psi_t(x_t)\psi_{st}(x_s, x_t) \prod_{u \in N(s) \setminus t} M_{us}^*(x_s) \prod_{u \in N(t) \setminus s} M_{ut}^*(x_t). \end{aligned}$$




Figure 1: (a) Undirected tree (for Problem 4.1(a)). (b), (c) Some undirected graphs to triangulate (Problem 4.4).

Problem 4.2

Undirected trees and marginals: Let $G = (V, E)$ be an undirected graph. For each vertex $i \in V$, let μ_i be a strictly positive function such that $\sum_{x_i} \mu_i(x_i) = 1$. For each edge, let μ_{ij} be a strictly positive function such that $\sum_{x_i} \mu_{ij}(x_i, x_j) = \mu_j(x_j)$ for all x_j , and $\sum_{x_j} \mu_{ij}(x_i, x_j) = \mu_i(x_i)$ for all x_i . Given integers k_1, \dots, k_d , consider the function

$$\mathbb{Q}(x_1, \dots, x_d) = \prod_{i=1}^d [\mu_i(x_i)]^{k_i} \prod_{(i,j) \in E} \mu_{ij}(x_i, x_j).$$

Supposing that G is a tree, can you give choices of integers k_1, \dots, k_d for which \mathbb{Q} is a valid probability distribution? If so, prove the validity. (*Hint:* It may be easiest to first think about a Markov chain.)

Solution: We claim that we can choose $k_i = 1 - \text{deg}_i$ to make $\mathbb{Q} \equiv \mathbb{Q}_G$ a probability distribution, where deg_i is the degree of node i in the tree G . We prove this claim by induction on the number of vertices d in the graph. When $d = 1$ the graph G only consists of a single node with degree 0, and $\mathbb{Q}_G(x_1) = \mu_1(x_1)$ is a probability distribution because $\sum_{x_1} \mu_1(x_1) = 1$. Now assume the claim is true for all trees with d vertices, and let $G = (V, E)$ be

a tree with $|V| = d + 1$ vertices. We want to prove that

$$\mathbb{Q}_G(x_1, \dots, x_d, x_{d+1}) = \prod_{i=1}^{d+1} \mu_i(x_i)^{1-\deg_i} \prod_{(i,j) \in E} \mu_{ij}(x_i, x_j)$$

is a probability distribution. Without loss of generality assume node $d + 1$ is a leaf in G , and its only neighbor is node d . Since $\mu_{d+1}(x_{d+1})$ appears with a power of $1 - \deg_{d+1} = 0$, the only factor involving x_{d+1} in \mathbb{Q}_G is $\mu_{d,d+1}(x_d, x_{d+1})$. Therefore, summing \mathbb{Q}_G over x_{d+1} gives us

$$\begin{aligned} \mathbb{Q}'_G(x_1, \dots, x_d) &= \sum_{x_{d+1}} \mathbb{Q}_G(x_1, \dots, x_d, x_{d+1}) \\ &= \left[\sum_{x_{d+1}} \mu_{d,d+1}(x_d, x_{d+1}) \right] \prod_{i=1}^d \mu_i(x_i)^{1-\deg_i} \prod_{\substack{(i,j) \in E \\ (i,j) \neq (d,d+1)}} \mu_{ij}(x_i, x_j) \\ &= \left(\prod_{i=1}^{d-1} \mu_i(x_i)^{1-\deg_i} \right) \mu_d(x_d)^{2-\deg_d} \prod_{(i,j) \in E \setminus \{(d,d+1)\}} \mu_{ij}(x_i, x_j). \end{aligned}$$

Let $G' = (V', E')$ be the tree that we get after removing vertex $d + 1$ from G , where $V' = V \setminus \{d + 1\}$ and $E' = E \setminus \{(d, d + 1)\}$, and note that the degrees of the nodes in G' now satisfy $\deg'_i = \deg_i$ for $1 \leq i \leq d - 1$, and $\deg'_d = \deg_d - 1$, since node d lost the edge $(d, d + 1)$. Now observe that we can write

$$\mathbb{Q}'_G(x_1, \dots, x_d) = \prod_{i=1}^d \mu_i(x_i)^{1-\deg'_i} \prod_{(i,j) \in E'} \mu_{ij}(x_i, x_j) = \mathbb{Q}_{G'}(x_1, \dots, x_d).$$

Therefore, using the inductive hypothesis that $\mathbb{Q}_{G'}$ is a probability distribution, we conclude that

$$\sum_{x_1, \dots, x_{d+1}} \mathbb{Q}_G(x_1, \dots, x_{d+1}) = \sum_{x_1, \dots, x_d} \mathbb{Q}_{G'}(x_1, \dots, x_d) = 1,$$

as desired.

Problem 4.3

Max-marginals: For each $i \in \{1, 2, \dots, d\}$, we define the max-marginal at node i via

$$q_i(x_i) = \max_{x_j, j \neq i} \mathbb{P}(x_1, x_2, \dots, x_d).$$

When x_i takes on m -states, then q_i specifies a vector of m numbers. This is the analog of an ordinary marginal distribution, with the summation replaced by maximization.

Suppose that $\arg \max_{x_i} q_i(x_i) = \{x_i^*\}$ for each node (i.e., the maximum is uniquely attained at x_i^*). Show that $(x_1^*, x_2^*, \dots, x_d^*)$ is the unique global optimum, meaning that $\mathbb{P}(x_1^*, \dots, x_d^*) > \mathbb{P}(x_1, \dots, x_d)$ for all $x \neq x^*$.

Solution: Let $x = (x_1, \dots, x_d) \in \arg \max_y \mathbb{P}(y)$. Then for each $1 \leq i \leq d$, on the one hand we have

$$\mathbb{P}(x) \leq \max_{y_j, j \neq i} \mathbb{P}(y_1, \dots, y_{i-1}, x_i, y_{i+1}, \dots, y_d) = q(x_i),$$

and on the other hand, because $\mathbb{P}(x) = \max_y \mathbb{P}(y)$, we also have

$$\mathbb{P}(x) \geq \max_{y_j, j \neq i} \mathbb{P}(y_1, \dots, y_{i-1}, x_i^*, y_{i+1}, \dots, y_d) = q(x_i^*) = \max_{y_i} q(y_i).$$

This means $x_i \in \arg \max_{y_i} q(y_i)$, which implies $x_i = x_i^*$. Since this holds for all $1 \leq i \leq d$, we conclude that $x = x^*$.

Problem 4.4

Triangulation/JT: Consider the two graphs shown in panels (b) and (c) of Figure 1. For each graph, first form a triangulated version, and then construct a junction tree using the greedy algorithm. (You can simply implement each step of the greedy algorithm on paper to find a maximum weight spanning tree.)

Solution:

- (a) We triangulate the graph by running elimination algorithm on it with ordering 1, 2, 3, 4, 5, 8, 10, 6, 7, 9 and using the reconstituted graph as the triangulated version of the initial graph (in general, we want to use elimination ordering that results in introducing as few new edges as possible). We then identify the maximum cliques and build a clique graph where the nodes are the maximum cliques and the edge weights are the cardinality of the intersection of the cliques. Finally, we construct the junction tree by finding a maximum weight spanning tree in this clique graph. See Figure 2 for a sample junction tree. (Note: this is just one of many possible solutions).
- (b) The tree-building pipeline is the same as in previous part. We use elimination ordering 4, 5, 1, 2, 3, 6, 7, 8, which only introduces 1 new edge.

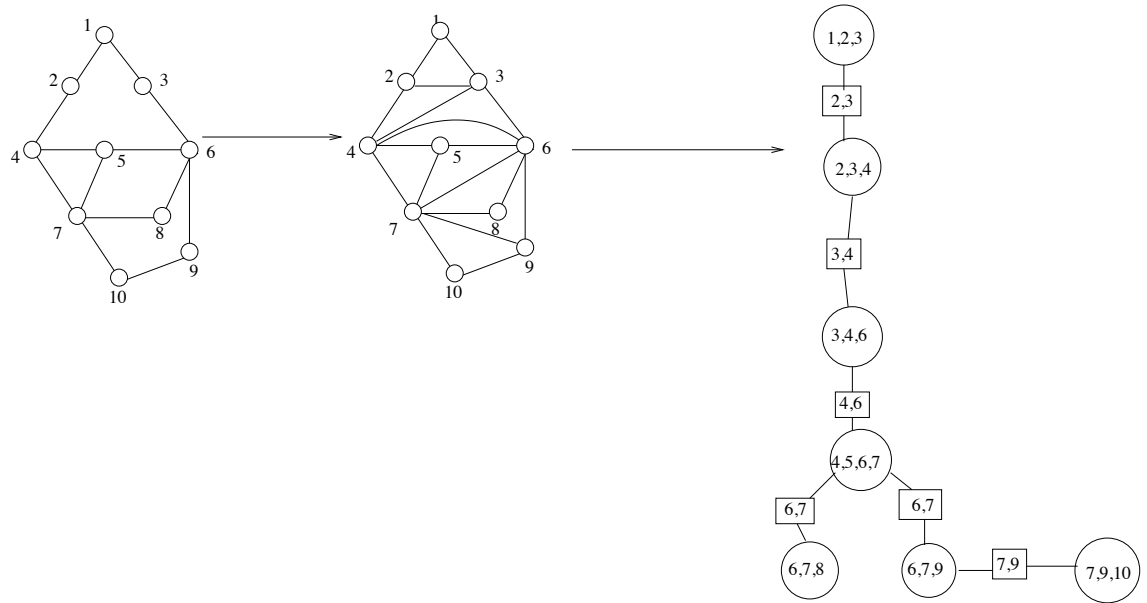


Figure 2: Problem 4.4(a): Initial graph, triangulated graph, and junction tree.