Probabilistic Graphical Models **Final Exam**	**Instructor:** **B. Nasihatkon**	K. N. TOOSI UNIVERSITY OF TECHNOLOGY
Name:	**ID:**	**Khordad 1396 - June 2017**

1. Variable Elimination (20 points)

Assume that we are to apply variable elimination to the following Bayesian network.

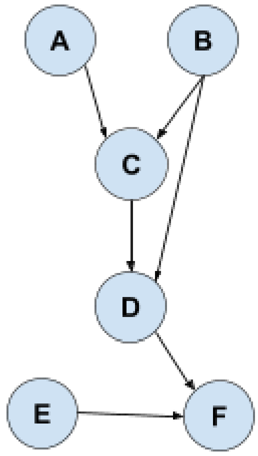

A) Draw the corresponding Markov network (4 points)

B) If you are to eliminate one variable from the network above, which variable(s) is (are) the best to start with? Which one is the worst? Why? (4 points)

C) To compute P(D | F), propose an optimal elimination ordering in terms of the algorithm efficiency. Why is this the best (or a best) order? (7 points)

D) Assume that we want to eliminate B, A, D, F and C in order. Draw the corresponding ***induced graph***. (5 points)

2. Junction Tree (22 points)

Consider the following junction tree (clique tree)

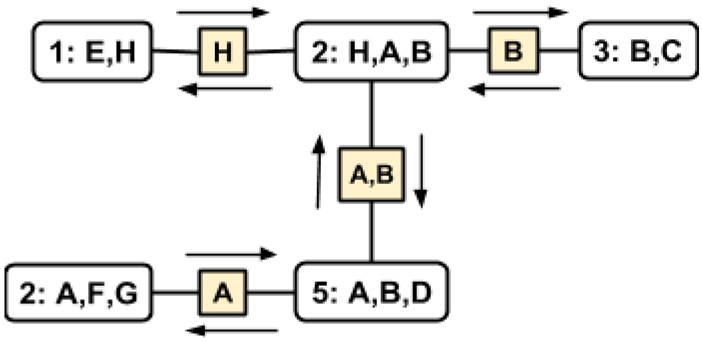

A) Write down the corresponding message beside each arrow. (2 points)

B) Draw a Markov network (MRF) corresponding to the above clique tree (4 points)

C) Write down the potential functions for the markov graph in part (B). Assume that there are **only binary and ternary** potentials and each cluster corresponds to **exactly one** potential function.For each potential function write down the corresponding cluster number (1,2,3,4 or 5). (3 points)

D) What is the minimum number of message computations needed for the Belief Propagation algorithm to converge to the right solution? Write one such ordering of messages. (6 points).

E) Assume that all variables are binary ($\in \{0,1\}$), $\phi_1(E,H) = exp(1(E = H))$, where 1(.) is the indicator function, $\phi_2(H,A,B) = exp(3\ A\ B\ H + 2A - BH)$, $\phi_3(B,C) = exp(2\ B\ C + B)$, and we are to perform **max-sum message passing** for **MAP** estimation. Derive $\delta_{1\to2}(H)$, $\delta_{3\to2}(B)$, **and then** $\delta_{2\to5}(A,B)$. Notice that the functions $\delta_{i\to j}$ are **max-sum** messages. You can either write a formula or a tabular representation. (7 points)

3. Random walk / MCMC (22 points)

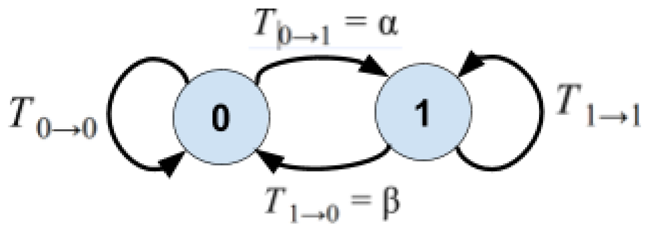

Consider a markov chain with the following transition model for a 1D distribution $P(X)$ with a binary variable $X \in \{0, 1\}$, in which $\alpha = T(0 \to 1)$ and $\beta = T(1 \to 0)$.

A) What are the values of $T(0 \to 0)$ and $T(1 \to 1)$ in terms of α and β. (1 point)

B) Assume that the transition probabilities α and β are given. Derive the corresponding stationary distribution $P^{\infty}(0) = \pi(0)$ and $P^{\infty}(1) = \pi(1)$ in terms of α and β. Write down the full derivations. (8 points)

C) Assume $\alpha = 0.3, \beta = 0.8$. What is the corresponding stationary distribution $P^{\infty}(X)$? (2 points)

D) Now, we want to solve the inverse problem. Assume that a special stationary distribution $P^{\infty}(X)$ is desired, that is, $P^{\infty}(0) = p$ and $P^{\infty}(1) = 1-p$ for a given p. We want to determine α and β in terms of p. Using the result from part (B), determine the ratio α/β in terms of p. Show that every solution $\alpha, \beta \in [0, 1]$ with this ratio is an answer to our problem, and therefore for a given stationary distribution the solution (α, β) is not unique. (Assume that $0 < p < 1$) (8 points)

E) Assume that we need to design a markov chain for which $P^{\infty}(0) = p = 0.4$. Obtain two different solutions (α, β) such that for the first one $\beta = 0.1$ and for the second one $\beta = 0.2$ (3 points).

F) ** Which of the two solutions in part (E) do you think gives a better random walk algorithm in terms of mixing more quickly? Give an intuitive explanation. Can you give an optimal solution (α, β) for part (E)? **(3+3 extra points)**

4. Parameter learning Bayesian Networks (16 points)

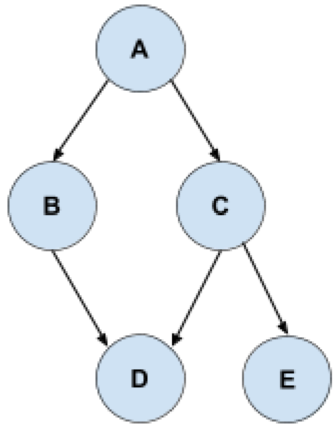

Consider the following Bayesian network with all-binary variables ($\in \{0,1\}$). Assume that the training data $X^1, X^2, \ldots, X^N$ is available, where $X^i = (a^i, b^i, c^i, d^i, e^i)$.

A) Write down the log-likelihood function in terms of the **logarithm of CPDs**. (6 points)

B) Assume that all the CPDs in this network are parameterized independently, except P(C | A) and P(E | C) which have shared parameters, i.e. share the same table, that is
P(C = x | A = y) = P(E = x| C = y).
Consider the following training data. Write the tabular representation of the Maximum Likelihood solution for each of the CPDs. (10 points)

	a^i	b^i	c^i	d^i	e^i
X^1	0	1	0	0	0
X^2	0	0	0	1	1
X^3	1	0	0	1	0
X^4	1	1	0	0	1
X^5	0	1	0	1	0
X^6	0	0	1	0	1

5. Parameter learning MRFs (20 points)

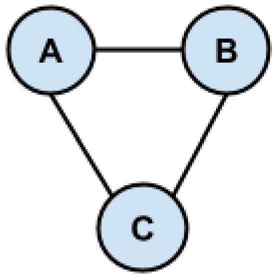

Consider the following MRF, on **binary** variables **A, B, C** $\in \{0, 1\}$ with joint distribution

$$P(A, B, C) = \frac{1}{Z} exp(w_1 \, 1(A = B) + w_2 \, 1(B = C) + w_3 \, 1(C = A) \,)$$

where **1(X = Y)** is equal to 1 if X = Y and zero otherwise.

A) Derive the partition function Z as a function of w_1, w_2, w_3. (4 points)

B) Consider the training data $X^1, X^2, \ldots, X^N$, where $X^i = (a^i, b^i, c^i)$. Write down the log-likelihood function in terms of the data (a^i, b^i, c^i) and the weights w_1, w_2, w_3. Simplify your answer as much as possible (4 points)

C) Derive the log-likelihood function for the following training data. Simplify your result as much as you can. (3 points)

	a^i	b^i	c^i
X^1	0	1	0
X^2	0	0	0
X^3	1	0	0
X^4	1	1	0
X^5	0	1	0
X^6	0	0	1

D) Which of the following assignments to w_1, w_2, w_3 better describes the data? Why? (3 points)

a) $w_1, w_2, w_3 = (2, 2, 1)$

b) $w_1, w_2, w_3 = (2, 1, 2)$

E) Take derivatives of the log-likelihood function of part (C) with respect to w_1, w_2 and w_3, and set them equal to zero. Let $a = e^{w_1}$, $b = e^{w_2}$, and $c = e^{w_3}$. Derive polynomial equations in terms of a, b, c for optimal w_1, w_2, w_3. (4 points)

F) Using the result of part (E) prove that for optimal parameters we have $w_1 = w_3$. (2 points)