Homework 5

MCMC, Learning

Milad Gholamrezanejad

Q1- MCMC

Consider a Markov chain with the following state transition graph for a 1D distribution, where the numbers by the edges are transition probabilities. Derive the corresponding stationary distribution $P_\infty(X = i) = \pi(i)$ for $i = 1, 2, 3, 4$. Write down the complete derivations.

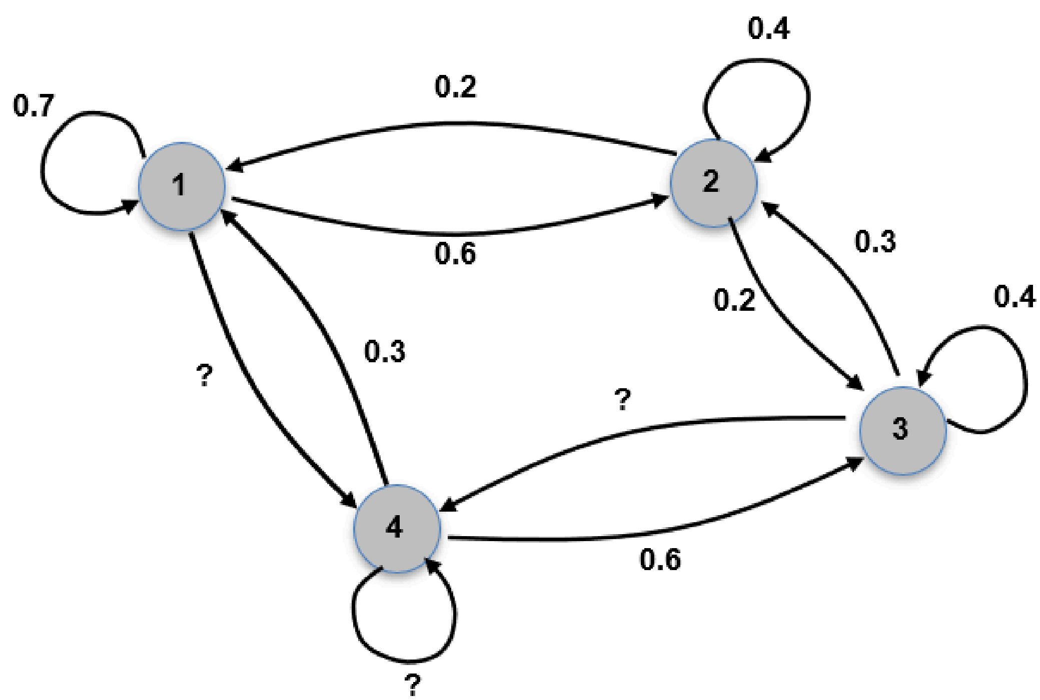

Q2 - Bayesian Network Parameter Learning

Consider the following Bayes Net on binary variables A, B, $C \in \{0, 1\}$, with CPDs defined as:

$P(A) = \alpha\, A + (1 - \alpha)\,(1 - A)$,
$P(B) = (1 + B) / 3$,
$P(C \mid A, B) = \gamma 1(C = A + B) + (1 - \gamma)\, 1(C \neq A + B)$.

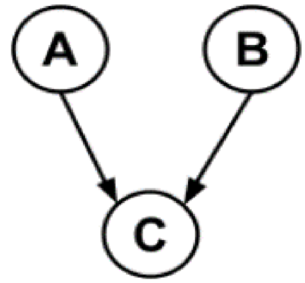

A) Write down the log-likelihood function in terms of α and γ and for the following data. Simplify your answer as much as possible. **Notice that some CPDs have not been parameterized by table representation.**

a^i	b^i	c^i
1	1	1
1	0	1
0	1	0
0	0	1
1	0	0
1	1	0

B) Write down the derivatives of the log-likelihood function for α, γ and find the optimal (maximum-likelihood) values of α and γ by setting the derivatives equal to zero.

Q3- Learning MRF

Consider the following MRF on binary variables $X, Y, Z \in \{0, 1\}$, with the log-linear joint distribution

$p(X, Y, Z) = \frac{1}{z} exp(\alpha\, 1(X = Y) + \beta\, YZ + \gamma 1(Z \neq X + Y))$

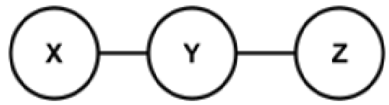

Your task is to obtain the optimal (maximum likelihood) values of the parameters α, β, and γ given the training data below. To do this you have to follow the steps below

a) Find the partition function Z as a function of α, β, and γ. Simplify as much as possible.

b) Write down the log-likelihood function using part (a) and the provided data.

c) Find the optimal values of α, β, and γ by differentiating with respect to each of these parameters and setting equal to zero.

Hint: *In case you need a change of variable, feel free to use u = exp(v).*

x^i	y^i	z^i
0	0	0
1	0	1
1	1	0
0	1	1
0	0	1