Probabilistic Artificial Intelligence
Final Exam
Jan 29, 2018

Time limit:	120 minutes
Number of pages:	15
Total points:	100

You can use the back of the pages if you run out of space. Collaboration on the exam is strictly forbidden. Please write your answers with a *pen*.

(1 point) Please fill in your student ID and full name (LASTNAME, FIRSTNAME) in capital letters.

Please leave the table below empty.

Problem	**Maximum points**	**Obtained**
1.	10	
2.	9	
3.	28	
4.	10	
5.	8	
6.	22	
7.	12	
Total	100	

1. Short Questions (10 points)

(10 points) For each of the statements below, state whether they are true or false. Each correct answer gives +1 point, each incorrect answer gives -1 point. You cannot get less than 0 points.

(a) In every Bayesian network there is at least one node with no parent.

☐ True ☐ False

(b) The correctness of the variable elimination algorithm depends on the choosing the right ordering.

☐ True ☐ False

(c) The variable elimination algorithm is used to perform exact inference.

☐ True ☐ False

(d) On a Bayesian network with n binary variables the variable elimination algorithm performs $\mathcal{O}(n^2)$ operations in the worst case.

☐ True ☐ False

(e) For polytrees, the maximum factor size during variable elimination is constant with respect to the number of variables, independent of the ordering chosen.

☐ True ☐ False

(f) For polytrees, variable elimination is generally more accurate than belief propagation.

☐ True ☐ False

(g) Inference algorithms for Bayesian networks can also be used to answer inference queries in HMMs.

☐ True ☐ False

(h) The domain of the observation variables in HMMs is equally sized to the domain of its state variables.

☐ True ☐ False

(i) An inefficient but technically correct way to compute the most probable explanation $(x_1^*, \ldots, x_T^*)$ in an HMM is to compute $x_t^* \in \operatorname{argmax}_x P(X_t = x \mid y_{1:T})$, for all $1 \leq t \leq T$.

☐ True ☐ False

(j) The only difference between HMMs and Kalman filters is that Kalman filters have a Gaussian sensor model.

☐ True ☐ False

2. Independence and d-separation (9 points)

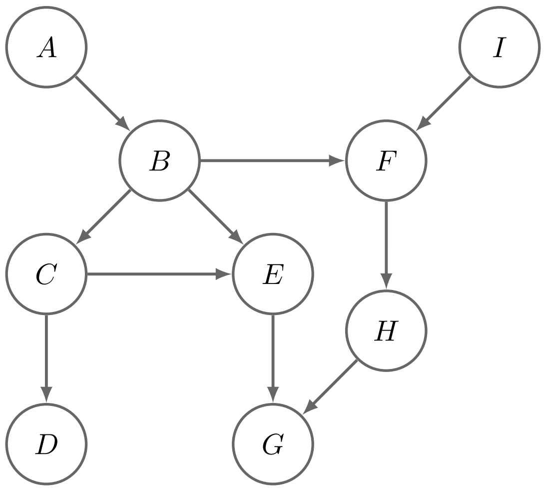

(9 points) For each of the statements below regarding the Bayesian network shown above, state whether they are true, false, or undecidable based on the information provided. Each correct answer gives +1.5 point, each incorrect answer gives -1.5 point. You cannot get less than 0 points.

(a) A and I are d-separated, that is, d-sep(A; I).

☐ True ☐ False ☐ Undecidable

(b) A and I are independent.

☐ True ☐ False ☐ Undecidable

(c) A and I are d-separated given G, that is, d-sep(A; $I \mid G$).

☐ True ☐ False ☐ Undecidable

(d) A and I are dependent given G.

☐ True ☐ False ☐ Undecidable

(e) There is a distribution that factorizes according to this network, such that A and I are dependent given B and H.

☐ True ☐ False ☐ Undecidable

(f) There is a distribution that factorizes according to this network, such that A and I are dependent given D and H.

☐ True ☐ False ☐ Undecidable

3. Roommates in Zurich (28 points)

Zurich is an expensive place to live, and so you and your friends share an apartment. As is not uncommon in these situations, you and your roommates have different ideas of what it means to live together. In particular, your roommates study easier subjects and have guests (G) over 50% of the nights. You, however, have to get up early in the morning and you always manage to get up (U) on time if your roommates had no guests over the night before. After a night with guests in the apartment, you sleep badly and get out of bed on time only with probability 0.5. On top of that, the apartment is dirty (D) if your roommates had a party and if you did not get out of bed on time to clean up after them. In all other cases, the apartment is clean. The reason you need to get up on time is that you need to catch the bus (B) going to ETH. You catch the bus with probability 0.8 if you got out of bed on time, and with probability 0.4 if you did not get up on time.

(7 points) (i) Draw the Bayesian network corresponding to the text above and write down the conditional probability distributions in terms of four boolean random variables G, U, D, B taking values True (T) or False (F).

(1 point) (ii) Write down the form of the joint probability distribution induced by the Bayesian network.

(3 points) (iii) Draw a factor graph with exactly two factor nodes representing the same distribution as above. Each factor may contain at most three variables. Define the two factors used.

(5 points) (iv) What is the probability of the apartment being dirty? Show the steps of your computation.

(6 points) (v) What is the probability of the apartment being dirty, given that you caught the bus? Show the steps of your computation.

(6 points) (vi) What is the probability that you caught the bus, given that the apartment is dirty? Show the steps of your computation.

4. Belief Propagation (10 points)

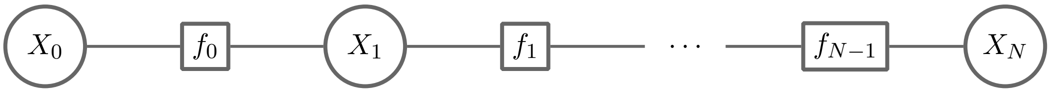

Consider the above factor graph corresponding to a chain Bayesian network. For $0 \leq i \leq N$, we assume that X_i takes values in $\{0, \ldots, M\}$ for some $M \in \mathbb{N}$. Recall that belief propagation computes $P(X_i = x)$, for any $i \in \{0, \ldots, N\}$ and any $x \in \{0, \ldots, M\}$.

(10 points) Suppose you would like to compute $P(X_i = x \mid X_0 \leq X_1 \leq \ldots \leq X_N)$. Show how this can be done using belief propagation on a modified factor graph. *(Hint: Show that the above factor graph can be transformed into a new factor graph of the same form, in which each factor f_i is replaced by a factor $\hat{f}_i$, so that computing $P(X_i = x)$ in the new graph gives the desired answer.)*

5. Learning Bayesian Networks (8 points)

Consider learning a Bayesian network of four variables X, Y, Z, W given a data set sampled from the joint distribution. The empirical pairwise mutual information has been computed as

$$\hat{I}(X;Y) = 0.32,\ \ \hat{I}(X;Z) = 0.38,\ \ \hat{I}(X;W) = 0.27,$$
$$\hat{I}(Y;Z) = 0.39,\ \ \hat{I}(Y;W) = 0.27,\ \ \hat{I}(Z;W) = 0.39.$$

Answer the following questions and briefly justify each answer.

(4 points) (i) Draw a Bayesian network that maximizes the likelihood of the observed data.

(4 points) (ii) Draw a tree-shaped Bayesian network that maximizes the likelihood of the observed data.

6. Server Scheduling (22 points)

Consider a server that serves two queues indexed by $i = 1, 2$. Queue i receives a new request at time-steps $2t + i$ ($t = 0, 1, 2, \ldots$). Each request remains in a queue for the duration of exactly one time step and expires afterwards, so that each queue can contain at most one request at a time.

Let q_t specify the queue the server serves at time t. Serving a queue proceeds as follows: if the queue is non-empty, then the server clears the request from that queue and earns a reward of 1; otherwise the server receives no reward.

At the end of each time step t, the server decides whether to remain at the current queue (i.e., $q_{t+1} = q_t$) or move to the other queue (i.e., $q_{t+1} \neq q_t$). If the server chooses to move, it incurs an additional cost of $C \geq 0$; otherwise, no cost is incured.

(8 points) (i) Draw the MDP described above, annotating the action-dependent transitions with transition probabilities and associated rewards.

(7 points) (ii) Let the discount factor be $\gamma = 0.5$ and the cost be $0 \leq C < 1/3$. Guess what the optimal policy is and solve the Bellman equations to compute the value function for that policy.

(7 points) (iii) Using the result of the previous question, show that the policy you guessed is indeed optimal.

7. Reinforcement Learning (12 points)

Consider the MDP represented by the following transition graph.

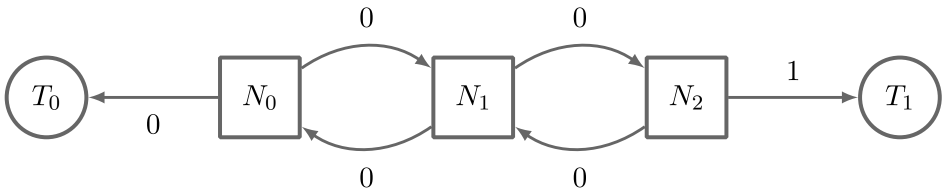

States T_0 and T_1 are terminal, while N_0, N_1, N_2 are non-terminal. (Remember that reaching a terminal state means staying in that state forever and gaining no further reward, i.e., signifies the end of an episode.) The action space consists of a single action a_0 that results in a uniformly random move to one of the neighboring states, for example, $P(X_{t+1} = T_0 \mid X_t = N_0, A = a_0) = P(X_{t+1} = N_1 \mid X_t = N_0, A = a_0) = 0.5$. Every move gives reward 0, except for the one from N_2 to T_1, which gives reward 1, as shown in the graph above.

(8 points) (i) Apply Q-learning to estimate the values of the Q-function of this MDP given the following episode traces. Show the update to the Q-values at each step.

Episode 1 $N_1 \to N_2 \to T_1$
Episode 2 $N_1 \to N_2 \to N_1 \to N_0 \to T_0$

The Q-values of all non-terminal states are initialized to 0.5, and the Q-values of the terminal states are known to be 0. (There is no need to show updates for terminal states.) Use discount factor $\gamma = 1$ and step size $\alpha = 0.5$.

(4 points) (ii) The following figures show the Q-values estimated from running the above algorithm for a large number of episodes, once using $\gamma = 0.4$, and another time using $\gamma = 0.7$. Match each figure to the corresponding discount factor and briefly explain your reasoning.

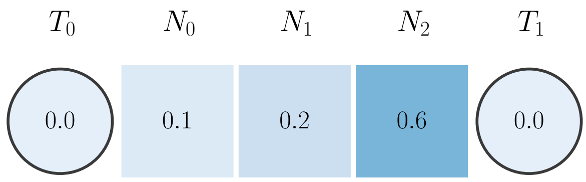

(a)

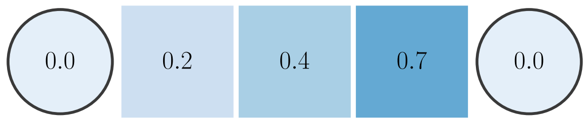

(b)