


1 Bayes' Nets

Bayesian networks represent our beliefs about the conditional independence relationships between the variables in a system. The joint distribution is factored into conditional distributions, which are represented by the connections between nodes.

1. What would a Bayes net of variables which are all independent of each other look like?
2. Consider the following Bayes net.



- (a) Which of the following statements is true? (Hint: d-separation!)
- $P(B | A) = P(B)$
 - $A \perp\!\!\!\perp B | D, F$
 - $A \perp\!\!\!\perp B | C$
 - $D \perp\!\!\!\perp E$
 - $P(D | C, E) = P(D | C)$
 - $D \perp\!\!\!\perp E | A, B$
 - $P(B | E) = P(B | G, E)$
 - $P(A | G) = P(A | B, G)$

2 Inference Conceptual Review

1. When would we want to use inference?

2. Suppose we are given binary random variables Q, H, E (query, hidden, evidence). We want to query $P(q | e)$.



(a) **Enumeration**

Perform inference on a joint distribution. Use the Bayes net above to break down joint into CPT factors.

Note: You may use a proportionality constant α in your answer.

(b) **Variable Elimination**

Rewrite your answer to enumeration by moving summations inwards as far as possible.


Note: You may use a proportionality constant α in your answer.

(c) Based on 2a and 2b, why is variable elimination more efficient than enumeration?

3 Inference

Realizing that students aren't particularly fond of reading the textbook, the 281 course staff have developed a software that automatically scans the textbook and outputs key points for each individual chapter. However, since the development of the software requires time and computational resources, the 281 staff decides to offer a free one month trial to students, after which a paid subscription is necessary to keep using the software. The following network and variables are used to represent the problem:

- $Discount(D)$: $+d$ if a discount is offered, $-d$ otherwise
- $Enjoys(E)$: $+e$ if a student enjoys the software, $-e$ otherwise
- $Cost(C)$: $+c$ if the software cost is < 20 , $-c$ otherwise
- $Recommends(R)$: $+s$ if the student recommends the software to a friend, $-s$ otherwise
- $Buys(B)$: $+b$ if the student buys a software subscription, $-b$ otherwise



1. How can we represent the probability that a student buys and recommends the software using the conditional probabilities at each node?
2. The staff has surveyed students and collected data on whether the students enjoyed the software or not. With this information, we want to perform inference on a joint distribution where the query variable is $Buys$ (B).
 - (a) How can we represent the probability expression in terms of conditional probabilities from the network?
 - (b) What are the hidden and evidence variable(s)?
3. Using the probability expression from the previous part, we want to compute the query B given evidence that the student enjoys the software. Assume the variable ordering is in alphabetical order.
 - (a) How many factors are there, and what are the dimensions of each factor?
 - (b) Run the variable elimination algorithm to eliminate repeated computations for the expression $P(B|+e)$.

- (c) How does the resulting expression change if the variable ordering is instead in reverse alphabetical order?





- (d) How do the two orderings compare with respect to time and space complexity?

- (e) Describe a heuristic that could be useful in determining a variable ordering to minimize the size of the largest factor.

4 Candy Sampling

- In the year 2020, the Oompa Loompas at Charlie’s Chocolate Factory have decided that they want to try a new automated way of sampling their candies for quality assurance. However, they have spilled chocolate sauce on their only copy of the user manual! Help out the Oompa Loompas by filling in the blanks below with the names of the four different types of sampling methods we’ve discussed in lecture, and then match each one to the corresponding image, probability distribution, and algorithm from the tables below.

Sampling Method Name	Image and Distribution (A-D)	Algorithm (1-4):

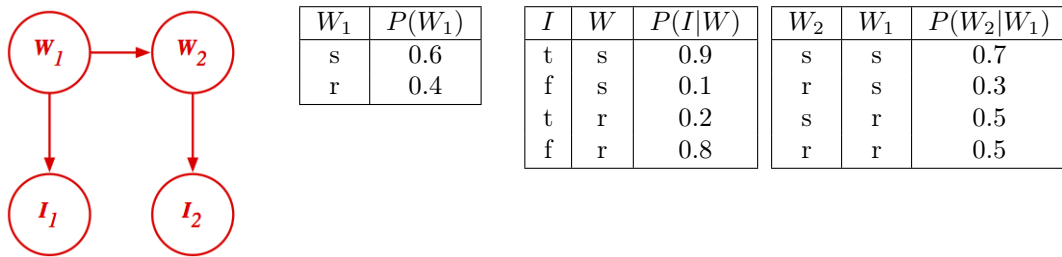
Name:	Images and Corresponding Probability Distributions:
(A)	 <p style="text-align: right;">$P(Q, E)$</p>
(B)	 <p style="text-align: right;">$P(Q e)$</p>
(C)	 <p style="text-align: right;">$P(Q, e)$</p>
(D)	 <p style="text-align: right;">$P(Q e)$</p>

Name:	Algorithms:
(1)	<pre> def likelihood_weighted_sample(evidence instantiation): w = 1.0 for i = 1,2,...n: if(X_i is an evidence variable): X_i = observation x_i for X_i w = w * P(x_i Parents(X_i)) else: Sample x_i from P(X_i Parents(X_i)) return ((x_1, x_2,...x_n), w) </pre>
(2)	<p>function _____(<i>bn</i>) returns an event sampled from the prior specified by <i>bn</i></p> <p>inputs: <i>bn</i>, a Bayesian network specifying joint distribution $\mathbf{P}(X_1, \dots, X_n)$</p> <p>$\mathbf{x} \leftarrow$ an event with <i>n</i> elements</p> <p>foreach variable X_i in X_1, \dots, X_n do</p> <p> $\mathbf{x}[i] \leftarrow$ a random sample from $\mathbf{P}(X_i \text{parents}(X_i))$</p> <p>return \mathbf{x}</p>
(3)	<p>function _____(<i>X, e, bn, N</i>) returns an estimate of $\mathbf{P}(X \mathbf{e})$</p> <p>local variables: <i>N</i>, a vector of counts for each value of <i>X</i>, initially zero</p> <p> <i>Z</i>, the nonevidence variables in <i>bn</i></p> <p> <i>x</i>, the current state of the network, initially copied from <i>e</i></p> <p>initialize <i>x</i> with random values for the variables in <i>Z</i></p> <p>for $j = 1$ to <i>N</i> do</p> <p> for each Z_i in <i>Z</i> do</p> <p> set the value of Z_i in <i>x</i> by sampling from $\mathbf{P}(Z_i mb(Z_i))$</p> <p> $\mathbf{N}[x] \leftarrow \mathbf{N}[x] + 1$ where <i>x</i> is the value of <i>X</i> in <i>x</i></p> <p>return NORMALIZE(<i>N</i>)</p>
(4)	<p>function _____(<i>X, e, bn, N</i>) returns an estimate of $\mathbf{P}(X \mathbf{e})$</p> <p>inputs: <i>X</i>, the query variable</p> <p> <i>e</i>, observed values for variables <i>E</i></p> <p> <i>bn</i>, a Bayesian network</p> <p> <i>N</i>, the total number of samples to be generated</p> <p>local variables: <i>N</i>, a vector of counts for each value of <i>X</i>, initially zero</p> <p>for $j = 1$ to <i>N</i> do</p> <p> $\mathbf{x} \leftarrow$ PRIOR-SAMPLE(<i>bn</i>)</p> <p> if <i>x</i> is consistent with <i>e</i> then</p> <p> $\mathbf{N}[x] \leftarrow \mathbf{N}[x] + 1$ where <i>x</i> is the value of <i>X</i> in <i>x</i></p> <p>return NORMALIZE(<i>N</i>)</p>

5 Sampling Practice

1. Compared to other sampling methods (rejection, likelihood weighting, Gibbs), what kind of information can prior sampling not use (that other methods can)?
2. How does rejection sampling work on a high level, and what is its biggest/immediate weakness?

The diagram below describes a person’s ice-cream eating habits based on the weather. The nodes W_i stand for the weather on a day i , which can either be **s** (sunny) or **r** (rainy). The nodes I_i represent whether the person ate ice-cream on day i , which can either be **t** (true) or **f** (false).



Assume we generate the following six samples given the evidence $I_1 = t$ and $I_2 = f$ using **Likelihood Weighted Sampling**:

$$(W_1, I_1, W_2, I_2) = \langle s, t, r, f \rangle, \langle r, t, r, f \rangle, \langle s, t, r, f \rangle, \langle s, t, s, f \rangle, \langle s, t, s, f \rangle, \langle r, t, s, f \rangle$$

Using these samples, we will complete the following table:

(W_1, I_1, W_2, I_2)	Count/N	w	Joint
s, t, s, f	2/6	0.09	0.03
s, t, r, f	/6		
r, t, s, f	/6		
r, t, r, f	/6		

1. What is the weight of the sample (s, t, r, f) above? Recall that the weight given to a sample in likelihood weighting is:

$$w = \prod_{\text{Evidence variables } e} P(e|\text{Parents}(e)).$$

2. What is the estimate of $P(s, t, r, f)$ given the samples?

3. Compute the rest of the entries in the table. Use the estimated joint probabilities to estimate $P(W_2 = r | I_1 = t, I_2 = f)$.

4. What is a weakness of likelihood weighing sampling? How does Gibbs sampling work, and how does it address this limitation?