Artificial Intelligence *CS5300, Spring 2009*

# HW8: Decision Diagrams

## 1 Sampling in the Alarm Network

Consider the Alarm Network; for instance, slide 2 of day18.

1. Suppose we ran some sampling algorithm in this domain and gathered the following samples:

```
~b  ~e  ~a  ~j  ~m
~b  ~e   a  ~j  ~m
~b  ~e  ~a   j  ~m
~b   e   a  ~j   m
~b  ~e  ~a  ~j  ~m
 b  ~e  ~a  ~j   m
~b   e   a   j  ~m
~b  ~e   a   j   m
```

   Please estimate the following quantities based on these sample:

   (a) $p(b)$
   (b) $p(b \mid a)$
   (c) $p(j \mid b)$
   (d) $p(e \mid b, \neg a)$

2. Do you think these samples are actually samples from the alarm network's probability distribution? Why or why not?

3. Suppose that we had the following samples with *likelihood weights*. Estimate the same quantities as before:

```
~b  ~e  ~a  ~j  ~m          1.3
~b  ~e   a  ~j  ~m          2.3
~b  ~e  ~a   j  ~m          0.2
 b  ~e  ~a  ~j   m          1.1
~b   e  ~a  ~j  ~m          0.5
 b  ~e  ~a   j  ~m          1.2
~b   e   a   j  ~m          5
~b  ~e   a   j   m          0.01
```

   Please estimate the following quantities based on these sample:

   (a) $p(b)$
   (b) $p(b \mid a)$
   (c) $p(j \mid b)$
   (d) $p(e \mid b, \neg a)$

4. Suppose I wanted to estimate $p(m \mid b)$. Would prior sampling work well in this case (why or why not)? What about rejection sampling? How many samples do you think you'd need to draw (order of magnitude) to get reasonable estimates for these two approaches?

## 2 Decisions in Class

Alice is back, and still trying to pass evil Clarence's class. This time she wants to model the class as a Bayes net. As before, there are two decisions she has to make. First, whether to study or not. Second, whether to sweet-talk Clarence or not. If she studies, she learns the material with probability 0.9. If she does not study, she learns the material with probability 0.01 (luck). If she sweet-talks Clarence, he ends up in a good mood with probability 0.7; if she does not, he is in a good mood with probability 0.5. Her final grade (whether she passes or not) depends on whether she learns the material *and* whether Clarence is in a good mood. If they're both true, she passes with probability 0.99. If she learns the material but Clarence is in a bad mood, she passes with probability 0.8. If she doesn't learn the material, but Clarence is in a good mood, she passes with probability 0.5. If she doesn't learns the material and Clarence is in a bad mood, she passes with probability 0.1.

1. Represent this situation as a Bayes net, complete with CPTs and decision nodes (appropriately shaped).
2. What is the value of information for knowing Clarence's mood?
3. (Extra credit for everyone!) Download the samiam program from Adnan Darwiche's website (at UCLA) and implement this in that set up (it's pretty easy). You probably will need to make the decision nodes just standard nodes. Do inference for all four possible action choices that Alice can make: what are the probabilities of passing for each. Please include a screenshot of the network in query mode.