Probabilistic Reasoning in AI - Assignment 3

Due Monday, February 18, 2008

1. [40 points] **Approximate vs exact inference**

   For this problem, you will experiment with the inference methods discussed so far in Matlab, using the Bayes nets toolbox, available at:
   http://www.cs.ubc.ca/ murphyk/Software/BNT/bnt.html

   You will be using the Insurance Bayes net, available in the Bayes net repository at:
   http://www.cs.huji.ac.il/labs/compbio/Repository/Datasets/insurance/insurance.htm

   (a) [10 points] Load the model into Matlab. Suppose that an adolescent who is a good student is coming to ensure an older economy car. Use the junction tree algorithm to compute the conditional probability of the PropCost variable. Record the answer you get, and, if possible, the computation time

   (b) [10 points] Repeat the experiment for loopy belief propagation.

   (c) [10 points] Use the likelihood weighting approximate inference algorithm to generate samples and compute these conditional probabilities approximately. In this case, generate increasingly large data sets. Plot a graph of the estimated probability as a function of the data set size. Add a line corresponding to the answer given by the junction tree algorithm. Comment in one-two sentences on the behavior you see in the graph.

   (d) [10 points] Do the same experiment but for Gibbs sampling. Produce the same sort of graph. Describe in a couple of sentences the behavior you observe.

2. [25 points] **Importance sampling**

   Suppose that x is a real-valued random variable with a Gaussian (normal) distribution with mean 1 and $\sigma = 1$. We want to determine the expected value of $f_1(x) = x$ and $f_2(x) = x^2$ from samples. But unfortunately we can only generate samples from a proposal distribution which is uniform between -1 and 10.

   (a) [10 points] Write a small routine, in the language of your choice, which generates samples , and then computes unnormalized and normalized importance sampling estimates of these expected values. Plot on the same graph the true values, and the estimated values and the number of samples increases.

   (b) [10 points] Now suppose that the proposal distribution is a Gaussian with mean 10 and standard deviation 1. What would you expect to see happening? Verify your prediction experimentally.

   (c) [5 points] Explain what will happen with the normalized and unnormalized importance sampling estimator as samples are drawn from similar, and respectively very different proposal distributions.

3. [10 points] **Markov chain**

   Consider the Markov chain shown below. Compute the n-step transition probability matrix:

$$p_{ij}^{(n)} = p(s_n = j | s_0 = i), \forall i, j \in \{0, 1, 2\}$$

   Does the chain have a steady-state distribution? If it does, compute it. Otherwise, justify why not.

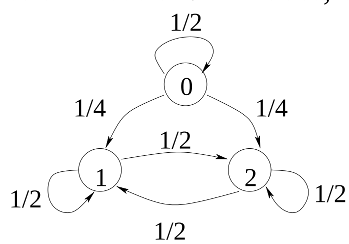

4. [25 points] **Gibbs sampling**

   Consider the Bayes net shown below:

$$X \to Z \leftarrow Y$$

   (a) [10 points] Assume X and Y are uniformly distributed, and Z is the deterministic exclusive or of X and Y. Show that Gibbs sampling on this structure with evidence $Z = 1$ will estimate $P(X = 1 | Z = 1)$ as either 1 or 0.

   (b) [15 points] What happens if we make Z a slightly noisy exclusive or? E.g. Z is the exclusive or of X and Y with probability $1 - q$ and is chosen uniformly randomly with probability q. Explain this behavior in terms of what happens with the random walk generated by Gibbs sampling.