Probabilistic Foundations of Artificial Intelligence
Problem Set 4

Oct 24, 2014

1. Belief propagation

In this exercise, you will implement the belief propagation algorithm for performing inference in Bayesian networks. As you have seen in the class lectures, the algorithm is based on converting the Bayesian network to a factor graph and then passing messages between variable and factor nodes of that graph until convergence.

You are provided some skeleton Python code in the `.zip` file accompanying this document. Take the following steps for this exercise.

1. Install the Python dependencies listed in `README.txt`, if your system does not already satisfy them. After that, you should be able to run `demo.py` and produce some plots, albeit wrong ones for now.

2. Implement the missing code in `bprop.py` marked with `TODO`. In particular, you have to fill in parts of the two functions that are responsible for sending messages from variable to factor nodes and vice versa, as well as parts of the function that returns the resulting marginal distribution of a variable node after message passing has terminated.

3. If your implementation is correct, you should get correct results for the naive Bayes model of the demo file that represents the coin flipping network of exercise 2 in Problem Set 2.

4. Now, set up the full-fledged earthquake network, whose structure was introduced in Problem Set 3 and is shown again in Figure 1. Here is the story behind this network:

   While Fred is commuting to work, he receives a phone call from his neighbor saying that the burglar alarm in Fred's house is ringing. Upon hearing this, Fred immediately turns around to get back and check his home. A few minutes on his way back, however, he hears on the radio that there was an earthquake near his home earlier that day. Relieved by the news, he turns around again and continues his way to work.

   To build up the conditional probability tables (CPTs) for the network of Figure 1 you may make the following assumptions about the variables involved:

   - All variables in the network are binary.
   - As can be seen from the network structure, burglaries and earthquakes are assumed to be independent. Furthermore, each of them is assumed to occur with probability 0.1%.
   - The alarm is triggered in the following ways: (1) When a burglar enters the house, the alarm will ring 99% of the time; (2) when an earthquake occurs, there will be a false alarm 1% of the time; (3) the alarm might go off due to other causes

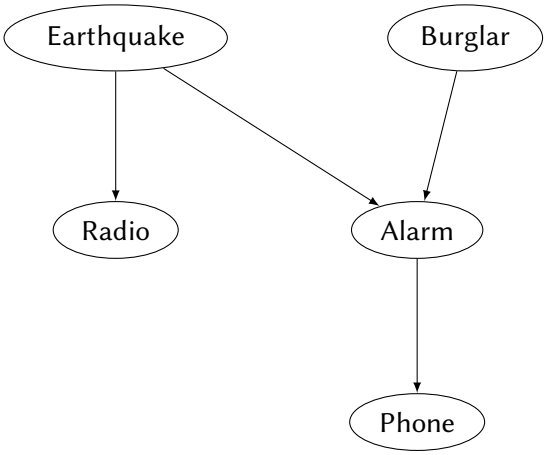

Figure 1: The earthquake network to be implemented.

(wind, rain, etc.) 0.1% of the time. These three types of causes are assumed to be independent of each other.

- The neighbor is assumed to call only when the alarm is ringing, but only does so 70% of the time when it is actually ringing.
- The radio is assumed to never falsely report an earthquake, but it might fail to report an earthquake that actually happened 50% of the time. (This includes the times that Fred fails to listen to the announcement.)

5. After having set up the network and its CPTs, answer the following questions using your belief propagation implementation:

   (a) Before Fred gets the neighbor's call, what is the probability of a burglary having occurred? What is the probability of an earthquake having occurred?

   (b) How do these probabilities change after Fred receives the neighbor's phonecall?

   (c) How do these probabilities change after Fred listens to the news on the radio?