CS 289 Knowledge representation and reasoning
Fall 1995 Stuart Russell **Project 3**

Stochastic simulation in belief nets

In this project, you will extend the IDEAL belief net package to include an inference algorithm that uses stochastic simulation. The IDEAL package is available as a tar file from the CS 289 WWW page. The WWW page also contains instructions for loading IDEAL.

The domain will be oil drilling. Specifically, imagine you are an oil industry analyst trying to predict whether LilOilCo will go bankrupt. This depends mainly on how much profit they make from their single oil well, which comes on stream in 1996. Relevant variables include the 1996 price of oil, the total operating cost (fixed costs for drilling and variable costs for extraction), and the total recoverable reserves. The principal influence on oil price is the peace process in the Middle East, which may or may not break down. The only specific data available is a single geologist's estimate of the recoverable reserves.

1. Construct a network of discrete variables, using appropriate discretizations (these can be very crude). Enter the network and associated conditional probabilities into IDEAL. Use IDEAL to calculate the probability of bankruptcy given a specific value for estimated reserves.

2. Write an algorithm to generate complete sample cases from the network in the absence of evidence, according to the distribution specified by the network.

3. Implement logic sampling and use it to calculate the probability of bankruptcy given the estimated reserves. Graph the absolute error as a function of the number of samples (average over several complete runs).

4. Implement likelihood weighting and repeat the experiment. Discuss the relative performance of the algorithms. How would this change if your discretizations were much finer?

5. Suppose you were to use continuous variables instead of discretization.

   (a) Describe the conditional density functions you would use in the network (be specific).

   (b) Explain how the likelihood weighting algorithm would change—including sampling, evidence likelihoods, and calculation of the desired probability estimates. What happens if the query variable is a continuous variable?

   (c) Sketch what the posterior distribution of profit should look like, given the estimated reserves and your specific network.