

e corresponds to the pre-defined evidence you set prior to running the network through AIS. AIS also employs two heuristics which adapt the conditional probability tables for unlikely evidence and extremely low probabilities by setting the parents of unlikely evidence nodes to a uniform distribution and by raising low probabilities to a set threshold. These two heuristics allow the posterior probabilities be learned properly despite very low prior probabilities. Because the heuristic initializations and the learned importance function allow AIS to learn the true probabilities better than traditional sampling, we currently use it as a fitness function for evaluating learned structures of BBNs based on different orderings.

K2 algorithm

The K2 algorithm, developed by Cooper and Herskovits [CH92], learns the structure of a BBN based on a collection of training examples and a specific ordering of the nodes. It is a greedy-search algorithm and only considers nodes previous to the current node when building the list of parents (thus, the order is very important). To evaluate the output of K2, the LS algorithm is used to compute each node's probabilities. We then compute root-mean-square-error (RMSE) between the probabilities learned from the samples and the probabilities computed by some inference method with the K2-produced network. RMSE values closest to 0.0 are the best outcomes of the network. We are using other means of judging the fitness of the network as well: adaptive importance sampling and forward sampling. Both are approximates to LS, which will be a little less accurate, but (theoretically) much faster for large networks.

K2 Genetic Algorithm Wrapper.

Finding the optimal ordering is *NP*-hard, so we have designed a genetic algorithm to avoid exhaustively searching all possibilities. RMSE values of the various orderings are used as the fitness. Each chromosome in our genetic algorithm corresponds to an ordering used in K2. We use an order-crossover operator as well as an order swap operator for mutation. The order-crossover's purpose is to maintain the order of the nodes while still swapping out with another chromosome. There are other methods of crossover that we will implement and test in order to improve the performance of our genetic algorithm. We have utilized *GAJIT* as the main genetic algorithm driver. *GAJIT* is an elitist GA. It allows the user to specify the cull rate (the percentage of chromosomes that make the cut to the next generation). It also allows the user to add self-designed crossover and mutation operators. There are some issues within *GAJIT* that we will attempt to address in the future, most notably the crowding affect. When selecting two chromosomes to crossover, *GAJIT* does not care if a chromosome crossovers with itself. Applying the order-crossover technique on two identical chromosomes will yield two identical children, thus filling the population with clones, making it difficult to get out of a local

optimum at times. We have developed a job-farm that allows several computers to work on the same generation simultaneously. The job farm utilizes TCP to communicate and is tolerant to lag or dropped connections. Because we use java, we can employ several machines, regardless of their platform. The server manages the genetic algorithm. When a client connects, the server sends important session parameters. Once the client is initialized, the server simply sends an ordering to the client. The client then runs the K2 algorithm on the ordering, runs the specified fitness function on the learned network, and finally sends back to the server the final fitness. The job farm is also capable of running an exhaustive permutation run on smaller networks. We have begun experiments with the K2 GA wrapper by comparing the fitness values among the gold standard Asia and Alarm networks, the canonical ordering (or any topologically sorted ordering), and the best ordering obtained through the wrapper. Our initial findings are very encouraging; the GA wrapper usually produces a network that differs from the gold standard by a small percentage of graph errors (missing, added, or reversed edges).

Genetic algorithm for finding the MPE.

Finding the Most Probable Explanation (MPE) of a network is very useful, although it is *NP*-hard. We have employed a GA to speed up the process by 'stacking the deck' as each generation progresses. A chromosome represents node value instantiations for all nodes in the network. Using *GAJIT* [Fa00] as our main GA driver, we have developed a Markov-blanket crossover and chromosome (sparse state coding) rotation as a mutation operator. The Markov blanket of radius n consists of a node's children, mates, parents, and the node itself, recursively computed n times for all nodes. We use a default radius of 1 in our GA. Although it is not guaranteed to find the MPE, the GA performs well given a sufficient number of generations and population size.

References

- [CD00] Cheng, J. and Druzel, M. J. 2000. AIS-BN: An adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research (JAIR)*, 13:155-188
- [CH92] Cooper, G. F. and Herskovits, E, 1992.. A Bayesian Method for the Induction of Probabilistic Networks from Data. *Machine Learning*, 9(4):309-347.
- [LS88] Lauritzen, S. L. and Spiegelhalter, D. J, 1988. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B* 50.