International Society for Bayesian Analysis, 9^{th} World Meeting,
Hamilton Island, Australia, 2008.

ON VARIABLE ELIMINATION IN DISCRETE BAYESIAN NETWORK INFERENCE

C.J. Butz[1⋆], W. Yan[1], P. Lingras[2], K. Konkel[1] and Y. Yao[1]

[1] University of Regina, Regina, Canada
[2] Saint Mary's University, Halifax, Canada
⋆ butz@cs.uregina.ca

We are interested in proving that *variable elimination* (VE) in discrete Bayesian networks always yields a clearly structured *conditional probability table* (CPT) rather than a potential as universally stated. A *Bayesian network* consists of a directed acyclic graph and a corresponding set of CPTs. Based on the conditional independencies holding in the directed acyclic graph, the product of the CPTs is a discrete joint probability distribution.

A fundamental step in processing queries is to remove variables from the factorization. VE removes a variable v by the following two steps:

(i) compute the product of all distributions involving v.

(ii) marginalize v out of the distribution produced in (i).

While it is always written that the VE algorithm produces a potential, we argue that VE output is, in fact, a CPT.

The difficult part in our work is showing that the product in step (i) is a CPT. To do so, we present an argument covering all possible situations that can arise during Bayesian network inference. Having established that the product in step (i) of VE always is a CPT, it is straightforward to show that step (ii) produces a CPT.

There are both practical and theoretical advantages to this study. It is widely acknowledged that understanding inference algorithms in Bayesian networks is an arduous task. In particular, probabilistic reasoning literature is not meant for readers looking for a way into the field. The use of potentials is perhaps the biggest hindrance to comprehension of probabilistic inference in Bayesian networks. Potentials do not have clear physical interpretation as they are not well-defined probability distributions. On the contrary, it is written that CPTs are easily interpretable.

Therefore, by describing variable elimination in Bayesian networks in terms of CPTs rather than potentials, our work has pedagogical value.

On the practical side, we were the first to argue that Bayesian network inference would be implemented in a relational database management system. In 2007, Corrada Bravo and Ramakrishnan empirically demonstrated in the SIGMOD conference that VE inference could be performed faster by utilizing database query optimization techniques. On the other hand, we would like to emphasize that the probabilistic relations in the above experiments are unstructured. That is, the probabilistic relations are viewed as being potentials. Our work outlined here establishes that there is a specific structure to the numbers in the probabilistic relations during VE inference. This CPT structure can be exploited during multiplication and marginalization for even faster inference.

We have already implemented a discrete Bayesian network shell that performs VE inference taking advantage of CPT structure. We have empirically tested our system using the seventy-seven large Bayesian networks given in the programming inference competition held at the 2006 Uncertainty in Artificial Intelligence conference. Similar to four of the five teams entered in the competition, our prototype system only finished on some of the networks. Of the fifty-eight networks we finished on, we obtained the correct result each and every time without exception.