Queries Outside a Clique

February 8, 2008

Consider a query $P(\boldsymbol{Y} \mid \boldsymbol{e})$ where the variables $\boldsymbol{Y}$ are not present together in a single clique. One naive approach is to construct a clique tree where we force one of the cliques to contain $\boldsymbol{Y}$. However, this approach forces us to tailor our clique tree to different queries, negating many of its advantages. An alternative approach is to perform variable elimination over a calibrated clique tree.

Example 0.1: Consider a clique tree obtained from a chain structured network A—B—C—D, with an appropriate set of factors $\mathcal{F}$. Our clique tree in this case would have three cliques $C_1 = \{A, B\}, C_2 = \{B, C\}$, and $C_3 = \{C, D\}$. Assume that we have calibrated the clique tree, so that the clique potentials are equal to the marginals of $P_{\mathcal{F}}$ as in Corollary 11.2.7. Assume that we now want to compute the probability $P_{\mathcal{F}}(B, D)$. If the entire clique tree is calibrated, so is any (connected) subtree $\mathcal{T}'$. Letting $\mathcal{T}'$ consist of the two cliques $\boldsymbol{C}_2$ and $\boldsymbol{C}_3$, it follows from Theorem 11.2.11 that:

$$P_{\mathcal{F}}(B, C, D) = \pi_{\mathcal{T}'}.$$

By Eq. (8.6), we have that:

$$\begin{aligned} P_{\mathcal{F}}(B, D) &= \sum_C P_{\mathcal{F}}(B, C, D) \\ &= \sum_C \frac{\pi_2[B, C]\pi_3[C, D]}{\mu_{2,3}(C)} \\ &= \sum_C P_{\mathcal{F}}(B \mid C) P_{\mathcal{F}}(C, D), \end{aligned}$$

where the last equality follows from calibration. Each of these probability expressions corresponds to a clique potential or a potential divided by a message. We can now perform variable elimination using these factors, in the usual way. ∎

More generally, we can compute the joint probability $P_{\mathcal{F}}(\boldsymbol{Y})$ for an arbitrary subset $\boldsymbol{Y}$ by using the potentials and messages in a calibrated clique tree to define factors corresponding to conditional probabilities in $P_{\mathcal{F}}$, and then performing variable elimination over the resulting set of factors. The precise

Procedure CTree-Query (
  $\mathcal{T}$, // Clique tree over $\mathcal{F}$
  $\{\pi_i\}, \{\mu_{i,j}\}$, // Calibrated clique and sepset potentials for $\mathcal{T}$
  $\boldsymbol{Y}$ // A query
)
  Let $\mathcal{T}'$ be a subtree of $\mathcal{T}$ such that $\boldsymbol{Y} \subseteq \mathit{Scope}[\mathcal{T}']$
  Select a clique $\boldsymbol{C}_r \in \mathcal{T}'$ to be the root
  $\mathcal{F} \leftarrow \pi_r$
  for each $\boldsymbol{C}_i \in \mathcal{T}'$
    $\phi \leftarrow \frac{\pi_i}{\mu_{i,p_r(i)}}$
    $\mathcal{F} \leftarrow \mathcal{F} \cup \{\phi\}$
  $\boldsymbol{Z} \leftarrow \mathit{Scope}[\mathcal{T}'] - \boldsymbol{Y}$
  Let $\prec$ be some ordering over $\boldsymbol{Z}$
  return Sum-Product-Variable-Elimination($\mathcal{F}, \boldsymbol{Z}, \prec$)

Figure 1: Out-of-clique inference in clique tree

algorithm is shown in Figure 1. The savings over simple variable elimination arise because we do not have to perform inference over the entire clique tree, but only over a portion of the tree that contains the variables $\boldsymbol{Y}$ which constitute our query. In cases where we have a very large clique tree, the savings can be significant.