A Fully Polynomial Time Approximation Scheme for Updating Credal Networks of Bounded Treewidth and Number of Variable States

Denis D. Mauá, Cassio P. de Campos, Marco Zaffalon
{denis,cassio,zaffalon}@idsia.ch

IDSIA, Switzerland

Extensively Specified Credal Networks

Graphical model

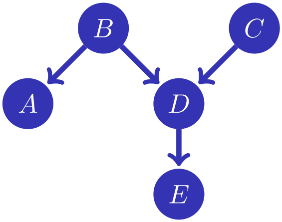

Extensively Specified Credal Networks

Discrete variables

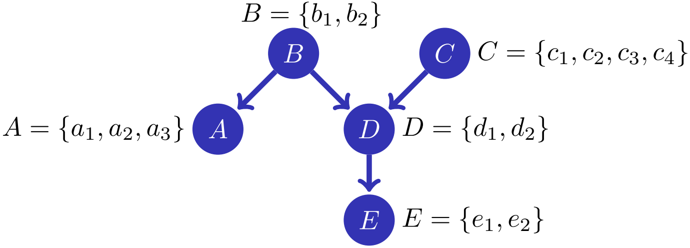

Extensively Specified Credal Networks

Local models

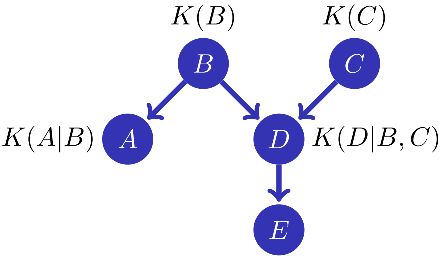

$$K(E|D) = \mathtt{H}\left\{ \begin{array}{|c|cc|}\hline & d_1 & d_2 \\ \hline e_1 & 0.2 & 0.6 \\ e_2 & 0.8 & 0.4 \\ \hline \end{array}, \begin{array}{|c|cc|}\hline & d_1 & d_2 \\ \hline e_1 & 0.3 & 0.2 \\ e_2 & 0.7 & 0.8 \\ \hline \end{array}, \begin{array}{|c|cc|}\hline & d_1 & d_2 \\ \hline e_1 & 0.9 & 0.6 \\ e_2 & 0.1 & 0.4 \\ \hline \end{array} \right\}$$

Extensively Specified Credal Networks

Inference

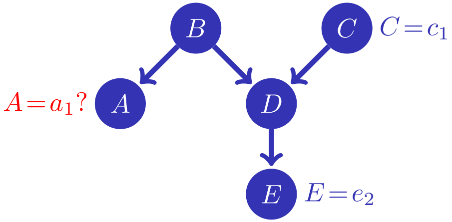

$$\max / \min\{P(a_1|c_1, e_2)\}$$

Belief Updating

Given evidence $E_1, \ldots, E_k = \mathbf{e}$ and query $Q = q$, let

$$K(Q, \mathbf{E}) = \Bigg\{ P(Q, \mathbf{E}) = \sum_{\mathbf{X} \setminus \{Q, \mathbf{E}\}} \prod_{i=1}^{n} P(X_i | \texttt{pa}(X_i)) : \\ P(X_i | \texttt{pa}(X_i) \in \texttt{ext}[K(X_i | \texttt{pa}(X_i))] \Bigg\}$$

The upper posterior probability is

$$\overline{P}(q|\mathbf{e}) = \max \left\{ \frac{P(q, \mathbf{e})}{P(q, \mathbf{e}) + P(\neg q, \mathbf{e})} : P(Q, \mathbf{E}) \in K(Q, \mathbf{E}) \right\}$$

Belief Updating as a Pareto Set

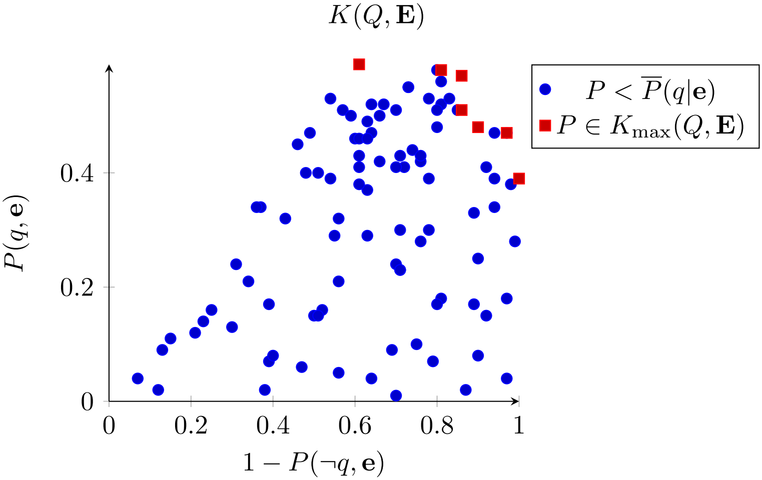

$$P(q|\mathbf{e}) = \frac{P(q, \mathbf{e})}{P(q, \mathbf{e}) + P(\neg q, \mathbf{e})}$$

Provably Good Approximate Inference

Goal is to obtain for any given $\epsilon > 1$ a solution $P(q|\mathbf{e})$ that is no worse than then optimum by a factor of ϵ:

$$P(q|\mathbf{e}) \geq \frac{\overline{P}(q|\mathbf{e})}{\epsilon}$$

- No polynomial-time algorithm exists if the variables can assume an arbitrary number of values, even in singly connected networks of bounded treewidth

Provably Good Approximate Inference as ϵ-Pareto Set

Relax Pareto dominance:

For $\alpha > 1$, $R(Q, \mathbf{E})$ α-dominates $P(Q, \mathbf{E})$ if

$$\alpha R(q, \mathbf{e}) \geq P(q, \mathbf{e})$$

and

$$R(\neg q, \mathbf{e}) \leq \alpha P(\neg q, \mathbf{e})$$

Hence,

$$\epsilon R(q|\mathbf{e}) \geq \alpha \frac{R(q, \mathbf{e})}{R(q, \mathbf{e}) + R(\neg q, \mathbf{e})} \geq \frac{P(q, \mathbf{e})}{P(q, \mathbf{e}) + P(\neg q, \mathbf{e})} = P(q|\mathbf{e})$$

Provably Good Approximate Inference as ϵ-Pareto Set

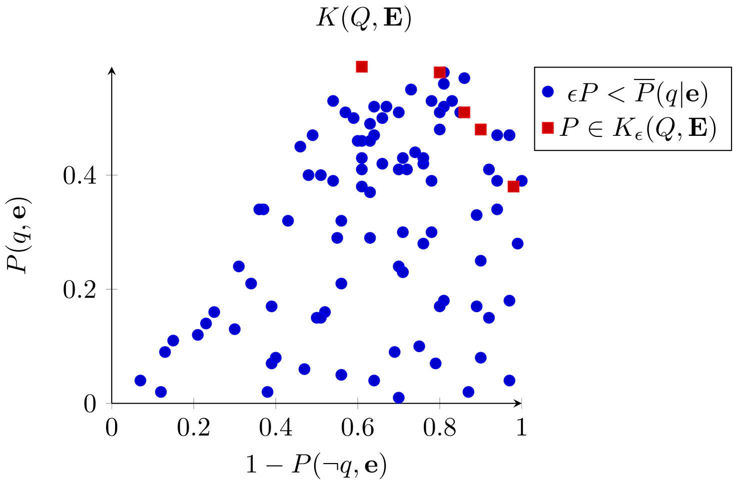

Provably Good Approximate Inference as ϵ-Pareto Set

Hence,

$$P_\epsilon(q|\mathbf{e}) = \max\left\{\frac{P(q,\mathbf{e})}{P(q,\mathbf{e}) + P(\neg q,\mathbf{e})} : P(Q,\mathbf{E}) \in K_\epsilon(Q,\mathbf{E})\right\}$$

is an ϵ-approximate solution.

- For networks of bounded treewidth and bounded number of states per variable, the number of elements in $K_\epsilon(Q,\mathbf{E})$ is bounded and can be computed in polynomial time by a message-passing scheme

Message-Passing Algorithm

- Operates over an ordered valuation algebra and a join tree whose nodes are associated to sets of valuations
- `prune` operation discards sub-optimal valuations

Exact Inference

- `prune` remove dominated valuations

$$\texttt{prune}\Psi = \{\phi \in \Psi : \nexists \psi \in \Psi \text{ such that } \phi \leq \psi \text{ and } \phi \neq \psi\}$$

Provably Good Approximate Inference

- Given $\epsilon > 1$, set $\alpha = 1 + \frac{\epsilon - 1}{4n}$
- `prune` remove α-equivalent valuations

$$\texttt{prune}\Psi = \{\phi \in \Psi : \nexists \psi \in \Psi \text{ such that } \phi \equiv_\alpha \psi\}$$

Provably Good Approximate Inference

- α-equivalence:

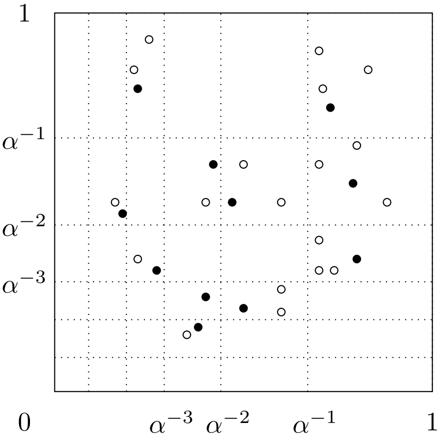

- It can be shown that the number of partitions is $O\left(\left[bn\frac{\alpha}{\alpha-1}\right]^{2k^{\omega}}\right)$

Experimental Results

Comparison against de Campos and Cozman's mixed integer linear programming algorithm:

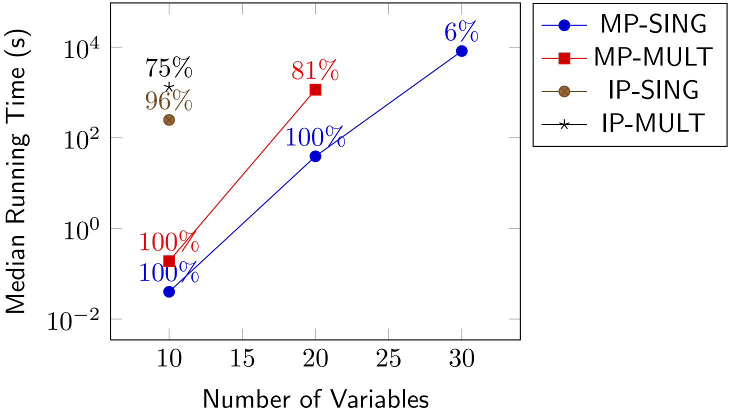

Experimental Results

Comparison against exact:

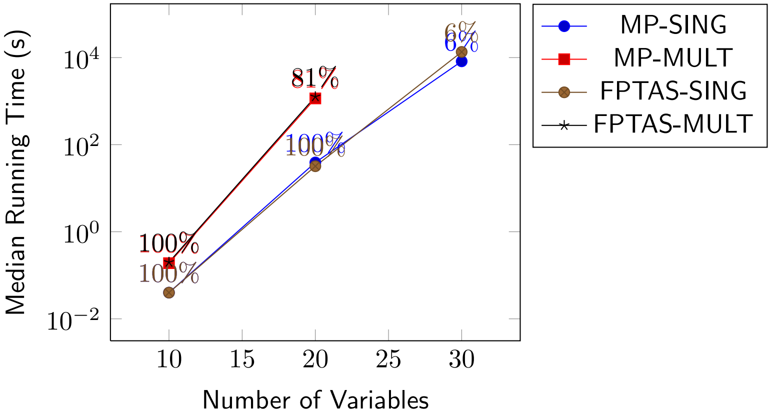

Experimental Results

Comparison against exact:

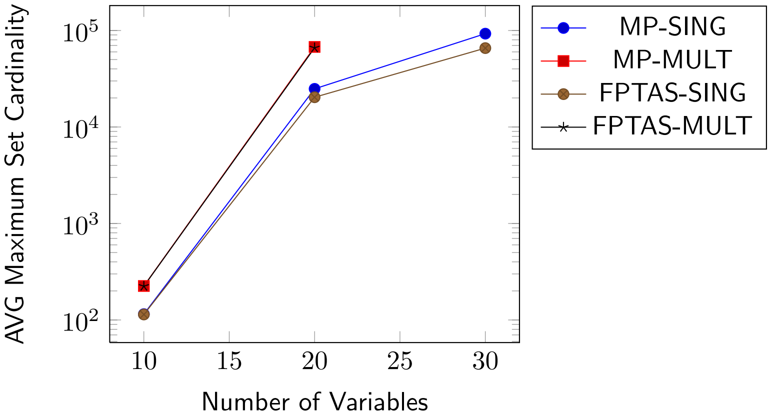

Questions?