
Efficient Principled Learning of Thin Junction Trees

Anton Chechetka Carlos Guestrin
Carnegie Mellon University

Abstract

We present the first truly polynomial algorithm for PAC-learning the structure of bounded-treewidth junction trees – an attractive subclass of probabilistic graphical models that permits both the compact representation of probability distributions and efficient exact inference. For a constant treewidth, our algorithm has polynomial time and sample complexity. If a junction tree with sufficiently strong intra-clique dependencies exists, we provide strong theoretical guarantees in terms of KL divergence of the result from the true distribution. We also present a lazy extension of our approach that leads to very significant speed ups in practice, and demonstrate the viability of our method empirically, on several real world datasets.

One of our key new theoretical insights is a method for bounding the conditional mutual information of arbitrarily large sets of variables with only polynomially many mutual information computations on fixed-size subsets of variables, if the underlying distribution can be approximated by a bounded-treewidth junction tree.

1 Introduction

In many applications, e.g., medical diagnosis or datacenter performance monitoring, **probabilistic inference** plays an important role: to decide on a patient’s treatment, it is useful to know the probability of various illnesses given the known symptoms. Thus, it is important to be able to represent probability distributions compactly and perform inference efficiently. Here, probabilistic graphical models (PGMs) have been successful as compact representations for probability distributions.

In order to use a PGM, one needs to define its structure and parameter values. Usually, we only have data (i.e., samples from a probability distribution), and learning the structure from data is thus a crucial task. For most formulations, the structure learning problem is NP-complete, *c.f.*, [10]. Most structure learning algorithms only guarantee that their output is a local optimum. One of the few notable exceptions is the work of Abbeel et al. [1], for learning structure of factor graphs, that provides probably approximately correct (PAC) learnability guarantees.

While PGMs can represent probability distributions compactly, exact inference in compact models, such as those of Abbeel et al., remains intractable [7]. An attractive solution is to use junction trees (JTs) of limited treewidth – a subclass of PGMs that permits efficient exact inference. For treewidth $k = 1$ (trees), the most likely (MLE) structure of a junction tree can be learned efficiently using the Chow-Liu algorithm [6], but the representational power of trees is often insufficient. We address the problem of learning JTs for fixed treewidth $k > 1$. Learning the most likely such JT is NP-complete [10]. While there are algorithms with global guarantees for learning fixed-treewidth JTs [10, 13], there has been no polynomial algorithm with PAC guarantees. The guarantee of [10] is in terms of the difference in log-likelihood of the MLE JT and the model where all variables are independent: the result is guaranteed to achieve at least a constant fraction of that difference. The constant does not improve as the amount of data increases, so it does not imply PAC learnability. The algorithm of [13] has PAC guarantees, but its complexity is exponential. In contrast, we provide a truly polynomial algorithm with PAC guarantees. The contributions of this paper are as follows:

- A theoretical result (Lemma 4) that upper bounds the conditional mutual information of arbitrarily large sets of random variables in polynomial time. In particular, we do not assume that an efficiently computable mutual information oracle exists.
- The first polynomial algorithm for PAC-learning the structure of limited-treewidth junction trees with strong intra-clique dependencies. We provide graceful degradation guarantees for distributions that are only approximately representable by JTs with fixed treewidth.

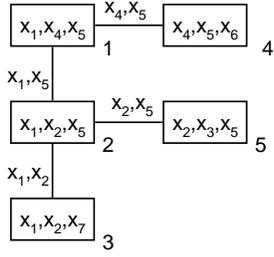


Figure 1: A junction tree. Rectangles denote cliques, separators are marked on the edges.

- A lazy heuristics that allows to make the algorithm practical.
- Empirical evidence of the viability of our approach on real-world datasets.

2 Bounded treewidth graphical models

In general, even to represent a probability distribution $P(V)$ over discrete variables¹ V we need space exponential in the size n of V . However, junction trees of limited treewidth allow compact representation *and* tractable exact inference. We briefly review junction trees (for details see [7]). Let $\mathbb{C} = \{C_1, \dots, C_m\}$ be a collection of subsets of V . Elements of \mathbb{C} are called **cliques**. Let T be a set of edges connecting pairs of cliques such that (T, \mathbb{C}) is a tree.

Definition 1. *Tree (T, \mathbb{C}) is a junction tree iff it satisfies the running intersection property (RIP): $\forall C_i, C_j \in \mathbb{C}$ and $\forall C_k$ on the (unique) simple path between C_i and C_j , $x \in C_i \cap C_j \Rightarrow x \in C_k$.*

A set $S_{ij} \equiv C_i \cap C_j$ is called the **separator** corresponding to an edge $(i-j)$ from T . The size of a largest clique in a junction tree minus one is called the **treewidth** of that tree. For example, in a junction tree in Fig. 1, variable x_2 is contained in both clique 3 and 5, so it has to be contained in clique 2, because 2 is on the simple path between 3 and 5. The largest clique in Fig. 1 has size 3, so the treewidth of that junction tree is 2.

A distribution $P(V)$ is representable using junction tree (T, \mathbb{C}) if instantiating all variables in a separator S_{ij} renders the variables on different sides of S_{ij} independent. Denote the fact that A is independent of B given C by $(A \perp B | C)$. Let \mathbb{C}_{ij}^i be cliques that can be reached from C_i in the (T, \mathbb{C}) without using edge $(i-j)$, and denote these reachable variables by $V_{ij}^i \equiv V_{ji}^i \equiv \bigcup_{C_k \in \mathbb{C}_{ij}^i} C_k \setminus S_{ij}$. For example, in Fig. 1, $S_{12} = \{x_1, x_5\}$, $V_{12}^1 = \{x_4, x_6\}$, $V_{12}^2 = \{x_2, x_3, x_7\}$.

Definition 2. $P(V)$ factors according to junction tree (T, \mathbb{C}) iff $\forall (i-j) \in T$, $(V_{ij}^i \perp V_{ij}^j | S_{ij})$.

If a distribution $P(V)$ factors according to some junction tree of treewidth k , we will say that $P(V)$ is **k -JT representable**. In this case, a **projection** $P_{(T, \mathbb{C})}$ of P on (T, \mathbb{C}) , defined as

$$P_{(T, \mathbb{C})} = \frac{\prod_{C_i \in \mathbb{C}} P(C_i)}{\prod_{(i-j) \in T} P(S_{ij})}, \quad (1)$$

is equal to P itself. For clarity, we will only consider **maximal** junction trees, where all separators have size k . If P is k -JT representable, it also factors according to some maximal JT of treewidth k .

In practice the notion of conditional independence is too strong. Instead, a natural relaxation is to require sets of variables to have low **conditional mutual information** I . Denote $H(A)$ the entropy of A , then $I(A, B | S) \equiv H(A | S) - H(A | BS)$ is nonnegative, and zero iff $(A \perp B | S)$. Intuitively, $I(A, B | S)$ shows how much new information about A can we extract from B if we already know S .

Definition 3. (T, \mathbb{C}) is an ε -junction tree for $P(V)$ iff $\forall (i-j) \in T$: $I(V_{ij}^i, V_{ij}^j | S_{ij}) \leq \varepsilon$.

¹Notation note: throughout the paper, we use small letters (x, y) to denote variables, capital letters (V, C) to denote sets of variables, and double-barred font (\mathbb{C}, \mathbb{D}) to denote sets of sets.

Algorithm 1: Naïve approach to structure learning

Input: V , oracle $I(\cdot, \cdot | \cdot)$, treewidth k , threshold δ

- 1 $\mathbb{L} \leftarrow \emptyset$; // \mathbb{L} is a set of “useful components”
- 2 **for** $S \subset V$ s.t. $|S| = k$ **do**
- 3 **for** $Q \subset V_S$ **do**
- 4 **if** $I(Q, V_{S \setminus Q} | S) \leq \delta$ **then**
- 5 $\mathbb{L} \leftarrow \mathbb{L} \cup (S, Q)$
- 6 **return** $FindConsistentTree(\mathbb{L})$

If there exists an ε -junction tree (T, \mathbb{C}) for $P(V)$, we will say that P is k -**JT** ε -**representable**. In this case, the Kullback-Leibler divergence of projection (1) of P on (T, \mathbb{C}) from P is bounded [13]:

$$KL(P, P_{(T, \mathbb{C})}) \leq n\varepsilon. \quad (2)$$

This bound means that if we have an ε -junction tree for $P(V)$, then instead of P we can use its tractable principled approximation $P_{(T, \mathbb{C})}$ for inference. In this paper, we address the problem of **learning structure** of such junction tree from data (samples from P).

3 Structure learning

In this paper, we address the following problem: given data, such as multiple temperature readings from sensors in a sensor network, we treat each datapoint as an instantiation of the random variables V and seek to find a good approximation of $P(V)$. We will assume that $P(V)$ is k -JT ε -representable for some ε and aim to find a $\hat{\varepsilon}$ -junction tree for P with the same treewidth k and with $\hat{\varepsilon}$ as small as possible. Note that the maximal treewidth k is considered to be a constant and not a part of problem input. The complexity of our approach is exponential in k .

Let us initially assume that we have an oracle $I(\cdot, \cdot | \cdot)$ that can compute the mutual information $I(A, B | C)$ exactly for any disjoint subsets $A, B, C \subset V$. This is a very strict requirement, which we address in the next section. Using the oracle I , a naïve approach would be to evaluate² $I(Q, V_{-QS} | S)$ for all possible $Q, S \subset V$ s.t. $|S| = k$ and record all pairs (S, Q) with $I(Q, V_{-QS} | S) \leq \varepsilon$ into a list \mathbb{L} . We will say that a junction tree (T, \mathbb{C}) is **consistent with a list** \mathbb{L} iff for every separator S_{ij} of (T, \mathbb{C}) it holds that $(S_{ij}, V_{ij}^i) \in \mathbb{L}$.

After \mathbb{L} is formed, any junction tree consistent with \mathbb{L} would be an ε -junction tree for $P(V)$. Such tree would be found by some *FindConsistentTree* procedure, implemented, e.g., using constraint satisfaction. Alg. 1 summarizes this idea. Algorithms that follow this outline, including ours, form a class of *constraint-based* approaches. These algorithms use mutual information tests to constrain the set of possible structures and return one that is consistent with the constraints. Unfortunately, using Alg. 1 directly is impractical because its complexity is exponential in the total number of variables n . In the following sections we discuss inefficiencies of Alg. 1 and present efficient solutions.

3.1 Global independence assertions from local tests

One can see two problems with the inner loop of Alg. 1 (lines 3-5). First, for each separator we need to call the oracle exponentially many times (2^{n-k-1} , once for every $Q \subset V_{-S}$). This drawback is addressed in the next section. Second, the mutual information oracle, $I(A, B | S)$, is called on subsets A and B of size $O(n)$. Unfortunately, the best known way of computing mutual information (and estimating I from data) has time and sample complexity exponential in $|A| + |B| + |S|$. Previous work has not addressed this problem. In particular, the approach of [13] has exponential complexity, in general, because it needs to estimate I for subsets of size $O(n)$. Our first new result states that we can limit ourselves to computing mutual information over small subsets of variables:

Lemma 4. *Let $P(V)$ be a k -JT ε -representable distribution. Let $S \subset V$, $A \subset V_{-S}$. If $\forall X \subseteq V_{-S}$ s.t. $|X| \leq k + 1$, it holds that $I(A \cap X, V_{-SA} \cap X | S) \leq \delta$, then $I(A, V_{-SA} | S) \leq n(\varepsilon + \delta)$.*

We can thus compute an upper bound on $I(A, V_{-SA} | S)$ using $O\left(\binom{n}{k}\right) \equiv O(n^k)$ (i.e., polynomially many) calls to the oracle $I(\cdot, \cdot | \cdot)$, and each call will involve at most $|S| + k + 1$ variables. Lemma 4 also bounds the quality of approximation of P by a projection on any junction tree (T, \mathbb{C}) :

Corollary 5. *If conditions of Lemma 4 hold for $P(V)$ with $S = S_{ij}$ and $A = V_{ij}^i$ for every separator S_{ij} of a junction tree (T, \mathbb{C}) , then (T, \mathbb{C}) is a $n(\varepsilon + \delta)$ -junction tree for $P(V)$.*

3.2 Partitioning algorithm for weak conditional independencies

Now that we have an efficient upper bound for $I(\cdot, \cdot | \cdot)$ oracle, let us turn to reducing the number of oracle calls by Alg. 1 from exponential (2^{n-k-1}) to polynomial. In [13], Narasimhan and Bilmes

²Notation note: for any sets A, B, C we will denote $A \setminus (B \cup C)$ as A_{-BC} to lighten the notation.

Algorithm 3: Efficient approach to structure learning

Input: V , oracle $I(\cdot, \cdot | \cdot)$, treewidth k , threshold ε , $\mathbb{L} = \emptyset$

- 1 **for** $S \subset V$ s.t. $|S| = k$ **do**
- 2 **for** $Q \in \text{LTCl}(V, S, I, \varepsilon, k + 2)$ **do**
- 3 $\mathbb{L} \leftarrow \mathbb{L} \cup (S, Q)$
- 4 **return** $\text{FindConsistentTreeDPGreedy}(\mathbb{L})$

Algorithm 4: FindConsistentTreeDPGreedy

Input: List \mathbb{L} of components (S, Q)

- 1 **for** $(S, Q) \in \mathbb{L}$ in the order of increasing $|Q|$ **do**
- 2 greedily check if (S, Q) is \mathbb{L} -decomposable
- 3 record the decomposition if it exists
- 4 **if** $\exists S : (S, V_S)$ is \mathbb{L} -decomposable **then**
- 5 **return** corresponding junction tree
- 6 **else return** no tree found

present an approximate solution to this problem, assuming that an efficient approximation of oracle $I(\cdot, \cdot | \cdot)$ exists. A key observation that they relied on is that the function $F_S(A) \equiv I(A, V_{-SA} | S)$ is **submodular**: $F_S(A) + F_S(B) \geq F_S(A \cup B) + F_S(A \cap B)$. Queyranne’s algorithm [14] allows the minimization of a submodular function F using $O(n^3)$ evaluations of F . [13] combines Queyranne’s algorithm with divide-and-conquer approach to partition V_S into conditionally independent subsets using $O(n^3)$ evaluations of $I(\cdot, \cdot | \cdot)$. However, since $I(\cdot, \cdot | \cdot)$ is computed for sets of size $O(n)$, complexity of their approach is still exponential in n , in general.

Our approach, called LTCl (Alg. 2), in contrast, has polynomial complexity for $q = O(1)$. We will show that $q = O(1)$ in our approach that uses LTCl as a subroutine. To gain intuition for LTCl, suppose there exists a ε -junction tree for $P(V)$, such that S is a separator and subsets B and C are on different sides of S in the junction tree. By definition, this means $I(B, C | S) \leq \varepsilon$. When we look at subset $A \equiv B \cup C$, the true partitioning is not known, but setting $\delta = \varepsilon$, we can test all possible $2^{|A|-1}$ ways to partition A into two subsets (X and A_X). If none of the possible partitionings have $I(X, A_X | S) \leq \varepsilon$, we can conclude that all variables in A are on the same side of separator S in any ε -junction tree that includes S as a separator. Notice also that

$$\forall X \subset A \quad I(X, A_X | S) > \delta \Leftrightarrow \min_{X \subset A} I(X, A_X | S) > \delta,$$

so we can use Queyranne’s algorithm to evaluate $I(\cdot, \cdot | \cdot)$ only $O(|A|^3)$ times instead of $2^{|A|-1}$ times for minimization by exhaustive search. LTCl initially assumes that every variable x forms its own partition $Q = \{x\}$. If a test shows that two variables x and y are on the same side of the separator, it follows that their container partitions $Q_1 \ni x, Q_2 \ni y$ cannot be separated by S , so LTCl merges Q_1 and Q_2 (line 3 of Alg. 2). This process is then repeated for larger sets of variables, of size up to q , until we converge to a set of partitions that are “almost independent” given S .

Proposition 6. *The time complexity of LTCl with $|S| = k$ is $O\left(\binom{n}{q} n J_{k+q}^{MI}\right) \equiv O\left(n^{q+1} J_{k+q}^{MI}\right)$, where J_{k+q}^{MI} is the time complexity of computing $I(A, B | C)$ for $|A| + |B| + |C| = k + q$.*

It is important that the partitioning algorithm returns partitions that are similar to connected components of V_{ij}^i of the true junction tree for $P(V)$. Formally, let us define two desirable properties. Suppose (T, \mathbb{C}) is an ε -junction tree for $P(V)$, and $\mathbb{Q}_{S_{ij}}$ is an output of the algorithm for separator S_{ij} and threshold δ . We will say that partitioning algorithm is **correct** iff for $\delta = \varepsilon$, $\forall Q \in \mathbb{Q}_{S_{ij}}$ either $Q \subseteq V_{ij}^i$ or $Q \subseteq V_{ij}^j$. A correct algorithm will never mistakenly put two variables on the same side of a separator. We will say that an algorithm is **α -weak** iff $\forall Q \in \mathbb{Q}_{S_{ij}} \quad I(Q, V_{-QS_{ij}} | S_{ij}) \leq \alpha$. For small α , an α -weak algorithm puts variables on different sides of a separator only if corresponding mutual information between those variables is not too large. Ideally, we want a correct and δ -weak algorithm; for $\delta = \varepsilon$ it would separate variables that are on different sides of S in a true junction tree, but not introduce any spurious independencies. LTCl, which we use instead of lines 3-5 in Alg. 1, satisfies the first requirement and a relaxed version of the second:

Lemma 7. *LTCl, for $q \geq k + 1$, is correct and $n(\varepsilon + (k - 1)\delta)$ -weak.*

3.3 Implementing FindConsistentTree using dynamic programming

A concrete form of *FindConsistentTree* procedure is the last step needed to make Alg. 1 practical. For *FindConsistentTree*, we adopt a dynamic programming approach from [2] that was also used in [13] for the same purpose. We briefly review the intuition; see [2] for details.

Consider a junction tree (T, \mathbb{C}) . Let S_{ij} be a separator in (T, \mathbb{C}) and \mathbb{C}_{ij}^i be the set of cliques reachable from C_i without using edge $(i - j)$. Denote T_{ij}^i the set of edges from T that connect

cliques from \mathbb{C}_{ij}^i . If (T, \mathbb{C}) is an ε -junction tree for $P(V)$, then $(\mathbb{C}_{ij}^i, T_{ij}^i)$ is an ε -junction tree for $P(V_{ij}^i \cup S_{ij})$. Moreover, the subtree $(\mathbb{C}_{ij}^i, T_{ij}^i)$ consists of a clique C_i and several sub-subtrees that are each connected to C_i . For example, in Fig. 1 the subtree over cliques 1,2,4,5 can be decomposed into clique 2 and two sub-subtrees: one including cliques $\{1,4\}$ and one with clique 5. The recursive structure suggests dynamic programming approach: given a component (S, Q) such that $I(Q, V_{-QS} \mid S) < \delta$, check if smaller subtrees can be put together to cover the variables of (S, Q) . Formally, we require the following property:

Definition 8. $(S, Q) \in \mathbb{L}$ is \mathbb{L} -decomposable iff $\exists \mathbb{D} = \cup_i \{(S_i, Q_i)\}, x \in Q$ s.t.

1. $\forall i (S_i, Q_i)$ is \mathbb{L} -decomposable and $\cup_{i=1}^m Q_i = Q \setminus \{x\}$;
2. $S_i \subset S \cup \{x\}$, i.e., each subcomponent can be connected directly to the clique (S, x) ;
3. $Q_i \cap Q_j = \emptyset$, ensuring the running intersection property within the subtree over $S \cup Q$.

The set $\{(S_1, Q_1), \dots, (S_m, Q_m)\}$ is called a **decomposition** of (S, Q) .

Unfortunately, checking whether a decomposition exists is equivalent to an NP-complete *exact set cover* problem because of the requirement $Q_i \cap Q_j = \emptyset$ in part 3 of Def. 8. Unfortunately, this challenging issue was not addressed by [13], where the same algorithm was used. To keep complexity polynomial, we use a simple greedy approach: for every $x \in Q_i$, starting with an empty candidate decomposition \mathbb{D} , add $(S_i, Q_i) \in \mathbb{L}$ to \mathbb{D} if the last two properties of Def. 8 hold for (S_i, Q_i) . If eventually Def. 8 holds, return the decomposition \mathbb{D} , otherwise return that no decomposition exists. We call the resulting procedure *FindConsistentTreeDPGreedy*.

Proposition 9. For separator size k , time complexity of *FindConsistentTreeDPGreedy* is $O(n^{k+2})$

Combining Alg. 2 and *FindConsistentTreeDPGreedy*, we arrive at Alg. 3. Overall complexity of Alg. 3 is dominated by Alg. 2 and is equal to $O(n^{2k+3} J_{2k+2}^{MI})$.

In general, *FindConsistentTreeDP* with greedy decomposition checks may miss a junction tree that is consistent with the list of components \mathbb{L} , but there is a class of distributions for which Alg. 3 is guaranteed to find a junction tree. Intuitively, we require that for every (S_{ij}, V_{ij}^i) from a ε -junction tree (T, \mathbb{C}) , Alg. 2 adds all the components from decomposition of (S_{ij}, V_{ij}^i) to \mathbb{L} and nothing else. This requirement is guaranteed for distributions where variables inside every clique of the junction tree are sufficiently strongly interdependent (have a certain level of mutual information):

Lemma 10. If \exists an ε -JT (T, \mathbb{C}) for $P(V)$ s.t. no two edges of T have the same separator, and for every separator S , clique $C \in \mathbb{C}$, $\min_{X \subset C-S} I(X, C-X_S \mid S) > (k+3)\varepsilon$ (we will call (T, \mathbb{C}) $(k+3)\varepsilon$ -strongly connected), then Alg. 3, called with $\delta = \varepsilon$, will output a $nk\varepsilon$ -JT for $P(V)$.

4 Sample complexity

So far we have assumed that a mutual information oracle $I(\cdot, \cdot \mid \cdot)$ exists for the distribution $P(V)$ and can be efficiently queried. In real life, however, one only has data (i.e., samples from $P(V)$) to work with. However, we can get a probabilistic estimate of $I(A, B \mid C)$, that has accuracy $\pm \Delta$ with probability $1 - \gamma$, using number of samples and computation time polynomial in $\frac{1}{\Delta}$ and $\log \frac{1}{\gamma}$:

Theorem 11. (Höffgen, [9]). The entropy of a probability distribution over $2k+2$ discrete variables with domain size R can be estimated with accuracy Δ with probability at least $(1 - \gamma)$ using $F(k, R, \Delta, \gamma) \equiv O\left(\frac{R^{4k+4}}{\Delta^2} \log^2\left(\frac{R^{2k+2}}{\Delta^2}\right) \log\left(\frac{R^{2k+2}}{\gamma}\right)\right)$ samples from P and the same amount of time.

If we employ this oracle in our algorithms, the performance guarantee becomes probabilistic:

Theorem 12. If there exists a $(k+3)(\varepsilon + 2\Delta)$ -strongly connected ε -junction tree for $P(V)$, then Alg. 3, called with $\delta = \varepsilon + \Delta$ and $\hat{I}(\cdot, \cdot, \cdot)$ based on Thm. 11, using $U \equiv F(k, R, \Delta, \frac{\gamma}{n^{2k+2}})$ samples and $O(n^{2k+3}U)$ time, will find a $kn(\varepsilon + 2\Delta)$ -junction tree for $P(V)$ with probability at least $(1 - \gamma)$.

Finally, if $P(V)$ is k -JT representable (i.e., $\varepsilon = 0$), and the corresponding junction tree is strongly connected, then we can let both Δ and γ go to zero and use Alg. 3 to find, with probability arbitrarily close to one, a junction tree that approximates P arbitrarily well in time polynomial in $\frac{1}{\Delta}$ and $\log \frac{1}{\gamma}$, i.e., the class of strongly connected k -junction trees is probably approximately correctly learnable³.

³A class \mathbb{P} of distributions is PAC learnable if for any $P \in \mathbb{P}$, $\delta > 0$, $\gamma > 0$ a learning algorithm will output $P' : KL(P, P') < \delta$ with probability $1 - \gamma$ in time polynomial in $\frac{1}{\delta}$ and $\log \frac{1}{\gamma}$.

Corollary 13. *If there exists an α -strongly connected junction tree for $P(V)$ with $\alpha > 0$, then for $\beta < \alpha n$, Alg. 3 will learn a β -junction tree for P with probability at least $1 - \gamma$ using $O\left(\frac{n^4}{\beta^2} \log^2 \frac{n}{\beta} \log \frac{n}{\gamma}\right)$ samples from $P(V)$ and $O\left(\frac{n^{2k+\tau}}{\beta^2} \log^2 \frac{n}{\beta} \log \frac{n}{\gamma}\right)$ computation time.*

5 Lazy evaluation of mutual information

Alg. 3 requires the value of threshold δ as an input. To get tighter quality guarantees, we need to choose the smallest δ for which Alg. 3 finds a junction tree. *A priori*, this value is not known, so we need a procedure to choose the optimal δ . A natural way to select δ is binary search. For discrete random variables with domain size R , for any $P(V), S, x$ it holds that $I(x, V_{-Sx} | S) \leq \log R$, so for any $\delta > \log R$ Alg. 3 is guaranteed to find a junction tree (with all cliques connected to the same separator). Thus, we can restrict binary search to range $\delta \in [0, \log R]$.

In binary search, for every value of δ , Alg. 2 checks the result of Queyranne’s algorithm minimizing $\min_{X \subset A} I(X, A_{-X} | S)$ for every $|S| = k, |A| \leq k+2$, which amounts to $O(n^{2k+2})$ complexity per value of δ . It is possible, however, to find the optimal δ while only checking $\min_{X \subset A} I(X, A_{-X} | S)$ for every S and A once over the course of the search process. Intuitively, think of the set of partitions \mathbb{Q}_S in Alg. 2 as a set of connected components of a graph with variables as vertices, and a hyper-edge connecting all variables from A whenever $\min_{X \subset A} I(X, A_{-X} | S) > \delta$. As δ increases, some of the hyper-edges disappear, and the number of connected components (or independent sets) may increase. More specifically, a graph \mathbb{Q}_S is maintained for each separator S . For all S, A add a hyper-edge connecting all variables in A annotated with $\text{strength}_S(A) \equiv \min_{X \subset A} I(X, A_{-X} | S)$ to \mathbb{Q}_S . Until $\text{FindConsistentTree}(\cup_S \mathbb{Q}_S)$ returns a tree, increase δ to be $\min_{S, A: \text{hyperedge}_S(A) \in \mathbb{Q}_S} \text{strength}_S(A)$ (i.e., strength of the weakest remaining hyper-edge), and remove $\text{hyperedge}_S(A)$ from \mathbb{Q}_S . Fig. 2(a) shows an example evolution of \mathbb{Q}_{x_4} for $k = 1$.

To further save computation time, we exploit two observations: First, if A is a subset of a connected component $Q \in \mathbb{Q}_S$, adding $\text{hyperedge}_S(A)$ to \mathbb{Q}_S will not change \mathbb{Q}_S . Thus, we do not test any hyper-edge A which is contained in a connected component. However, as δ increases, a component may become disconnected, because such an edge was not added. Therefore, we may have more components than we should (inducing incorrect independencies). This issue is addressed by our second insight: If we find a junction tree for a particular value of δ , we only need to recheck the components used in this tree. These insights lead to a simple, *lazy* procedure: If $\text{FindConsistentTree}$ returns a tree (T, \mathbb{C}) , we check the hyper-edges that intersect the components used to form (T, \mathbb{C}) . If none of these edges are added, then we can return (T, \mathbb{C}) for this value of δ . Otherwise, some of \mathbb{Q}_S have changed; we can iterate this procedure until we find a solution.

6 Evaluation

To evaluate our approach, we have applied it to two real-world (sensor network temperature [8] and San Francisco Bay area traffic [11]) and one artificial (samples from ALARM Bayesian network [4]) datasets. Our implementation, called LPACJT, uses lazy evaluations of $I(\cdot, \cdot | \cdot)$ from section 5. As baselines for comparison, we used a simple hill-climbing heuristic⁴, a combination of LPACJT with hill-climbing, where intermediate results returned by $\text{FindConsistentTree}$ were used as starting points for hill-climbing, Chow-Liu algorithm, and algorithms of [10] (denoted Karger-Srebro) and [17] (denoted OBS). All experiments were run on a Pentium D 3.4 GHz, with runtimes capped to 10 hours. The necessary entropies were cached in advance.

ALARM. This discrete-valued data was sampled from a known Bayesian network with treewidth 4. We learned models with treewidth 3 because of computational concerns. Fig. 2(b) shows the per-point log-likelihood of learned models on test data depending on the amount of training data. We see that on small training datasets both LPACJT finds better models than a basic hill-climbing approach, but worse than the OBS of [17] and Chow-Liu. The implementation of OBS was the only one to use regularization, so this outcome can be expected. We can also conclude that on this dataset our approach overfits than hill-climbing. For large enough training sets, LPACJT results achieve the likelihood of the true model, despite being limited to models with smaller treewidth. Chow-Liu performs much worse, since it is limited to models with treewidth 1. Fig. 2(c) shows an example of a structure found by LPACJT for ALARM data. LPACJT only missed 3 edges of the true model.

⁴Hill-climbing had 2 kinds of moves available: replace variable x with variable y in a connected sub-junction tree, or replace a leaf clique C_i with another clique $(C_i \setminus S_{ij}) \cup S_{mr}$ connected to a separator S_{mr} .

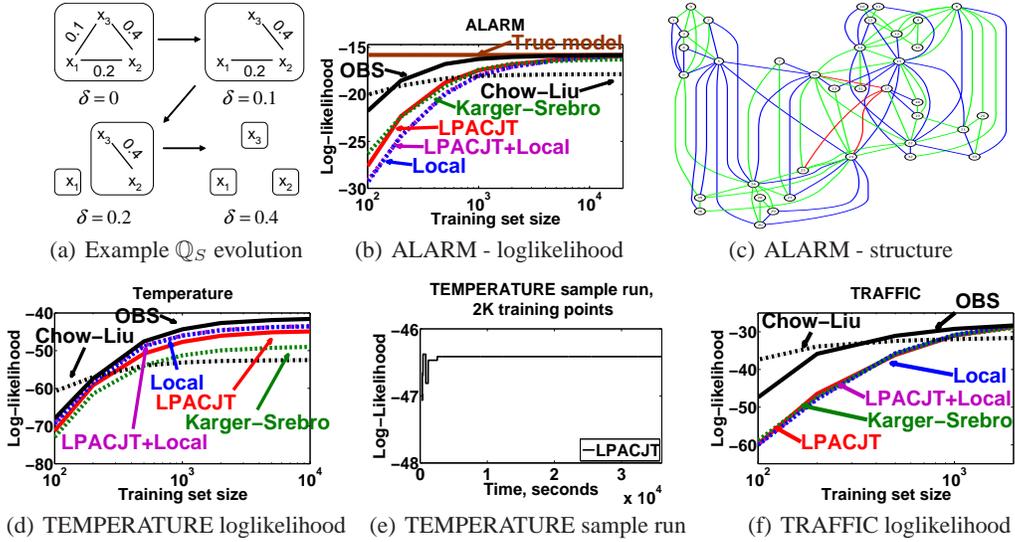


Figure 2: An example of evolution of \mathbb{Q}_S for section 5 (2(a)), one structure learned by LPACJT(2(c)), experimental results (2(b),2(d),2(f)), and an example evolution of the test set likelihood of the best found model (2(e)). In 2(c), nodes denote variables, edges connect variables that belong to the same clique, green edges belong to both true and learned models, blue edges belong only to the learned model, red - only to the true one.

TEMPERATURE. This data is from a 2-month deployment of 54 sensor nodes (15K data-points) [8]. Each variable was discretized into 4 bins and we learned models of treewidth 2. Since the locations of the sensor have an ∞ -like shape with two loops, the problem of learning a thin junction tree for this data is hard. In Fig. 2(d) one can see that LPACJT performs almost as good as hill-climbing-based approaches, and, on large training sets, much better than Karger-Srebro algorithm. Again, as expected, LPACJT outperforms Chow-Liu algorithm by a significant margin if there is enough data available, but overfits on the smallest training sets. Fig 2(e) shows the evolution of the test set likelihood of the best (highest training set likelihood) structure identified by LPACJT over time. The first structure was identified within 5 minutes, and the final result within 1 hour.

TRAFFIC. This dataset contains traffic flow information measured every 5 minutes in 8K locations in California for 1 month [11]. We selected 32 locations in San Francisco Bay area for the experiments, discretized traffic flow values into 4 bins and learned models of treewidth 3. All non-regularized algorithms, including LPACJT, give results of essentially the same quality.

7 Relation to prior work and conclusions

For a brief overview of the prior work, we refer the reader to Fig. 3. Most closely related to LPACJT are learning factor graphs of [1] and learning limited-treewidth Markov nets of [13, 10]. Unlike our approach, [1] does not guarantee low treewidth of the result, instead settling for compactness. [13, 10] guarantee low treewidth. However, [10] only guarantees that the difference of the log-likelihood of the result from the fully independent model is within a constant factor from the difference of the most likely JT: $LLH(\text{optimal}) - LLH(\text{indep.}) \leq 8^k k!^2 (LLH(\text{learned}) - LLH(\text{indep.}))$. [13] has exponential complexity. Our approach has polynomial complexity and quality guarantees that hold for strongly connected k -JT ε -representable distributions, while those of [13] only hold for $\varepsilon = 0$.

We have presented the first truly polynomial algorithm for learning junction trees with limited treewidth. Based on a new upper bound for conditional mutual information that can be computed using polynomial time and number of samples, our algorithm is guaranteed to find a junction tree that is close in KL divergence to the true distribution, for *strongly connected* k -JT ε -representable distributions. As a special case of these guarantees, we show PAC-learnability of strongly connected k -JT representable distributions. We believe that the new theoretical insights herein provide significant step in the understanding of structure learning in graphical models, and are useful for the analysis of other approaches to the problem. In addition to the theory, we have also demonstrated experimentally that these theoretical ideas are viable, and can, in the future, be used in the development of fast and effective structure learning heuristics.

approach	model class	guarantees	true distribution	samples	time	reference
score	tractable	local	any	any	poly [†]	[3, 5]
score	tree	global	any	any	$O(n^2)$	[6]
score	tree mixture	local	any	any	$O(n^2)$ [†]	[12]
score	compact	local	any	any	poly [†]	[17]
score	all	global	any	any	exp	[15]
score	tractable	const-factor	any	any	poly	[10]
constraint	compact	PAC ^o	positive	poly	poly	[1]
constraint	all	global	any	∞	poly(tests)	[16]
constraint	tractable	PAC	strong k -JT	exp [‡]	exp [‡]	[13]
constraint	tractable	PAC ^s	strong k -JT	poly	poly	this paper

Figure 3: Prior work. The majority of the literature can be subdivided into score-based [3, 5, 6, 12, 15, 10] and constraint-based [13, 16, 1] approaches. The former try to maximize some target function, usually regularized likelihood, while the latter perform conditional independence tests and restrict the set of candidate structures to those consistent with the results of the tests. *Tractable* means that the result is guaranteed to be of limited treewidth, *compact* - with limited connectivity of the graph. *Guarantees* column shows whether the result is a local or global optimum, whether there are PAC guarantees, or whether the difference of the log-likelihood of the result from the fully independent model is within a *const-factor* from the difference of the most likely JT. *True distribution* shows for what class of distributions the guarantees hold. † superscript means per-iteration complexity, poly - $O(n^{O(k)})$, exp[‡] - exponential in general, but poly for special cases. PAC^o and PAC^s mean PAC with (different) graceful degradation guarantees.

8 Acknowledgments

This work is supported in part by NSF grant IIS-0644225 and by the ONR under MURI N000140710747. C. Guestrin was also supported in part by an Alfred P. Sloan Fellowship. We thank Nathan Srebro for helpful discussions, and Josep Roure, Ajit Singh, CMU AUTON lab, Mark Teyssier, Daphne Koller, Percy Liang and Nathan Srebro for sharing their source code.

References

- [1] P. Abbeel, D. Koller, and A. Y. Ng. Learning factor graphs in polynomial time and sample complexity. *JMLR*, 7, 2006.
- [2] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic and Discrete Methods*, 8(2):277–284, 1987.
- [3] F. R. Bach and M. I. Jordan. Thin junction trees. In *NIPS*, 2002.
- [4] I. Beinlich, J. Suermondt, M. Chavez, and G. Cooper. The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. In *Euro. Conf. on AI in Medicine*, 1988.
- [5] A. Choi, H. Chan, and A. Darwiche. On Bayesian network approximation by edge deletion. In *UAI*, 2005.
- [6] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.
- [7] R. G. Cowell, P. A. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems (Information Science and Statistics)*. Springer, May 2003.
- [8] A. Deshpande, C. Guestrin, S. Madden, J. Hellerstein, and W. Hong. Model-driven data acquisition in sensor networks. In *VLDB*, 2004.
- [9] K. U. Höffgen. Learning and robust learning of product distributions. In *COLT*, 1993.
- [10] D. Karger and N. Srebro. Learning Markov networks: Maximum bounded tree-width graphs. SODA-01.
- [11] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. *UAI-05*.
- [12] M. Meilă and M. I. Jordan. Learning with mixtures of trees. *JMLR*, 1:1–48, 2001.
- [13] M. Narasimhan and J. Bilmes. PAC-learning bounded tree-width graphical models. In *UAI*, 2004.
- [14] M. Queyranne. Minimizing symmetric submodular functions. *Math. Programming*, 82(1):3–12, 1998.
- [15] A. Singh and A. Moore. Finding optimal Bayesian networks by dynamic programming. Technical Report CMU-CALD-05-106, Carnegie Mellon University, Center for Automated Learning and Discovery, 2005.
- [16] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT Press, 2001.
- [17] M. Teyssier and D. Koller. Ordering-based search: A simple and effective algorithm for learning Bayesian networks. In *UAI*, 2005.

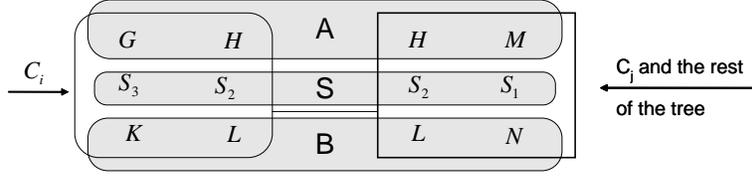


Figure 4: Illustration of notation in the proof of Lemma 4.

A Proofs of the theorems

To prove our results, we will make use of the following useful properties of conditional mutual information: for disjoint sets W, X, Y, Z and any probability distribution $P(WXYZ)$ the following holds:

- **Chain rule:** $I(W, XY | Z) = I(W, X | Z) + I(W, Y | XZ)$.
- **Nonnegativity:** $I(X, Y | Z) \geq 0$.
- **Monotonicity:** $X' \subseteq X, Y' \subseteq Y \Rightarrow I(X', Y' | Z) \leq I(X, Y | Z)$.

Proof of Lemma 4. We will prove slightly more general statement:

Lemma. *Let $P(V)$ be a probability distribution. Let A, B, S be a partitioning of V . Suppose there exists an ε -junction tree (T, \mathbb{C}) for $P(V)$ s.t. for every $C \in \mathbb{C}$ $|C \cap AB| \leq k + 1$. If $\forall X \subset AB$ s.t. $|X| \leq k + 1$ it holds that*

$$I(A \cap X, B \cap X | S) \leq \delta,$$

then

$$I(A, B | S) \leq |AB|(\varepsilon + \delta).$$

Note that in this formulation a clique C of a junction tree (T, \mathbb{C}) may have size greater than $k + 1$, up to $k + 1 + |S|$.

The proof is by induction. Base case: (T, \mathbb{C}) contains only one clique. Then, taking $X \equiv AB$, we have

$$I(A \cap AB, B \cap AB | S) \leq \delta \Leftrightarrow I(A, B | S) \leq \delta \Rightarrow I(A, B | S) \leq |AB|(\varepsilon + \delta)$$

(the last transition is valid, because $|AB| \geq 2$), so the theorem statement is true.

Now suppose (T, \mathbb{C}) contains more than one clique. Consider a leaf clique $C_i \in \mathbb{C}$ that is connected to only one other clique C_j . If a tree is not empty, a leaf clique always exists. Let S_{ij} be the corresponding separator. Define (see Fig. 4 for illustration)

$$G \equiv (C_i \setminus S_{ij}) \cap A, \quad H \equiv S_{ij} \cap A, \quad M \equiv A_{-GH}$$

$$K \equiv (C_i \setminus S_{ij}) \cap B, \quad L \equiv S_{ij} \cap B, \quad N \equiv B_{-KL}.$$

WLOG assume that both G and K are not empty. By the chain rule,

$$\begin{aligned} I(A, B | S) &= I(GHM, KLN | S) \\ &= I(HM, KLN | S) + I(G, KLN | SHM) \\ &= I(HM, LN | S) + I(HM, K | SLN) + I(G, KLN | SHM) \end{aligned} \quad (3)$$

Unroll the last two terms using the chain rule:

$$\begin{aligned}
I(G, KLN \mid SHM) &= I(G, MKLN \mid SH) - I(G, M \mid SH) \\
I(G, MKLN \mid SH) &= I(G, MN \mid KLSH) + I(G, KL \mid SH) \\
I(G, KL \mid SH) &= I(GH, KL \mid S) - I(H, KL \mid S) \\
I(G, MN \mid KLSH) &= I(GK, MN \mid LSH) - I(K, MN \mid LSH) \\
&\downarrow \\
I(G, KLN \mid SHM) &= I(GK, MN \mid LSH) - I(K, MN \mid LSH) + \\
&\quad + I(GH, KL \mid S) - I(H, KL \mid S) - I(G, M \mid SH) \\
&\text{and} \\
I(HM, K \mid SLN) &= I(HMN, K \mid SL) - I(N, K \mid SL) \\
I(HMN, K \mid SL) &= I(H, K \mid SL) + I(MN, K \mid HSL) \\
I(H, K \mid SL) &= I(H, KL \mid S) - I(H, L \mid S) \\
&\downarrow \\
I(HM, K \mid SLN) &= I(H, KL \mid S) - I(H, L \mid S) + I(MN, K \mid HSL) - I(N, K \mid SL)
\end{aligned}$$

and plug them back into (3):

$$\begin{aligned}
I(A, B \mid S) &= I(HM, LN \mid S) + I(HM, K \mid SLN) + I(G, KLN \mid SHM) \\
&= I(HM, LN \mid S) \\
&\quad + I(\mathbf{H}, \mathbf{KL} \mid \mathbf{S}) - I(H, L \mid S) + I(\mathbf{MN}, \mathbf{K} \mid \mathbf{HSL}) - I(N, K \mid SL) \\
&\quad + I(GK, MN \mid LSH) - I(\mathbf{K}, \mathbf{MN} \mid \mathbf{LSH}) + \\
&\quad + I(GH, KL \mid S) - I(\mathbf{H}, \mathbf{KL} \mid \mathbf{S}) - I(G, M \mid SH) \\
&= I(HM, LN \mid S) + I(GK, MN \mid LSH) + I(GH, KL \mid S) \\
&\quad - I(H, L \mid S) - I(N, K \mid SL) - I(G, M \mid SH) \\
&\leq I(HM, LN \mid S) + I(GK, MN \mid LSH) + I(GH, KL \mid S) \\
&\leq I(HM, LN \mid S) + I(GK, MN \mid LSH) + \delta \tag{4}
\end{aligned}$$

Let us show that $I(GK, MN \mid LSH) \leq \varepsilon$. Denote

$$S_1 = S \setminus C_i, \quad S_2 = S \cap S_{ij}, \quad S_3 = S \setminus S_1 S_2.$$

Note also that $S_2 LH \equiv S_{ij}$. Using chain rule, we have

$$\begin{aligned}
I(GK, MN \mid LSH) &= I(GK, MN \mid LS_1 S_2 S_3 H) \\
&= I(GK S_3, MN \mid LS_1 S_2 H) - I(S_3, MN \mid LS_1 S_2 H) \\
&= I(GK S_3, MNS_1 \mid LS_2 H) - I(GK S_3, S_1 \mid LS_2 H) - I(S_3, MN \mid LS_1 S_2 H) \\
&\leq I(GK S_3, MNS_1 \mid LS_2 H) \equiv I(C_i, V_{ij}^j \mid S_{ij}) \leq \varepsilon
\end{aligned}$$

Plugging this back into (4), we get

$$I(A, B \mid S) \leq I(HM, LN \mid S) + \varepsilon + \delta = I(A_{-G}, B_{-K} \mid S) + \varepsilon + \delta$$

the theorem. We now have to make an induction step. Consider a junction tree (T^*, \mathbb{C}^*) , which is the same as (T, \mathbb{C}) , but with C_i removed and with S_3 added to C_j :

$$\mathbb{C}^* = (\mathbb{C} \setminus \{C_i, C_j\}) \cup \{C_j \cup S_3\}.$$

Then (T^*, \mathbb{C}^*) is an ε -JT for $P(A_{-G} B_{-K} S)$. Indeed, consider any separator S_{ml} from (T^*, \mathbb{C}^*) . It is also a separator in (T, \mathbb{C}) . Denote V_{ml}^{m*} the set of all variables on one side of S_{ml} in (T^*, \mathbb{C}^*) , and V_{ml}^m - all variables on the same side, but in (T, \mathbb{C}) . We have $V_{ml}^{m*} \subseteq V_{ml}^m$ (equality if C_i is on the same side of S_{ml} as C_m , and strict subset if they are on different sides). Analogously, $V_{ml}^{l*} \subseteq V_{ml}^l$ and thus

$$I(V_{ml}^{m*}, V_{ml}^{l*} \mid S_{ml}) \leq I(V_{ml}^m, V_{ml}^l \mid S_{ml}) \leq \varepsilon$$

Therefore, (T^*, \mathbb{C}^*) is an ε -JT for $P(A_{-G} B_{-K} S)$, and for every $X \subset A_{-G} B_{-K}$, $|X| \leq k+1$ it holds that $X \subset AB$ and

$$I(X \cap A_{-G}, X \cap B_{-K} \mid S) = I(X \cap A, X \cap B \mid S) \leq \delta.$$

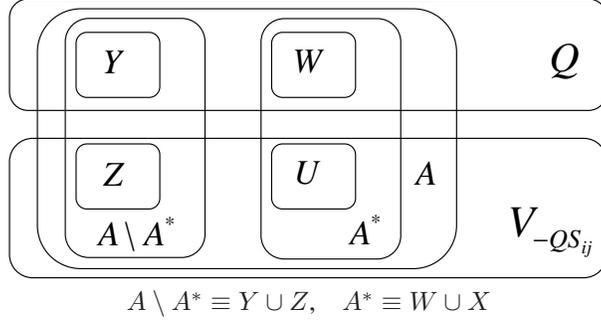


Figure 5: Illustration of notation in the proof of Lemma 7.

Finally, $\forall C_k^* \in \mathbb{C}^*$ it holds that

$$AB \cap C_k^* \equiv AB \cap C_k \Rightarrow |A_{-G}B_{-K} \cap C_k^*| \leq k + 1$$

so the induction conditions hold. By induction, then

$$I(A_{-G}, B_{-K} \mid S) \leq (|A_{-G}| + |B_{-K}|)(\epsilon + \delta) = (|AB| - |GK|)(\epsilon + \delta)$$

and since $|GK| \geq 1$,

$$I(A, B \mid S) = \epsilon + \delta + I(A_{-G}, B_{-K} \mid S) \leq |AB|(\epsilon + \delta)$$

□

Proof of Corollary 5. From Lemma 4, for every separator S_{ij} it holds that

$$I(V_{ij}^i, V \setminus (V_{ij}^i S_{ij}) \mid S_{ij}) \leq n(\epsilon + \delta)$$

and since in a connected junction tree $V \setminus (V_{ij}^i S_{ij}) \equiv V_{ij}^j$, by definition (T, \mathbb{C}) is an $n(\epsilon + \delta)$ -junction tree for $P(V)$. □

Proof of Proposition 6. The algorithm cycles through all subsets of size at most $k+1$ of V on line 2. This results in $\sum_{i=1}^{k+1} \binom{n}{i+1} = O(n^{k+1})$ iterations. For each iteration, it has to run Queyranne's algorithm to find

$$\min_X I(X, A_{-X} \mid S).$$

Queyranne's algorithm takes $O(|A|^3) = O(k^3) = \{k \text{ is a constant}\} = O(1)$ mutual information oracle calls to find this minimum. Thus the time spent on it for every subset A is $O(Q_{|A|+|S|}^{MI}) \leq O(Q_{2k+1}^{MI})$. Finally, for every A , finding the matching components $Q \in \mathbb{Q}_S, Q \cap A \neq \emptyset$ and merging them takes at most $O(n)$ time. □

Proof of Lemma 7. The proof of correctness is by induction. Initially, each variable x form $V_{-S_{ij}}$ forms its own partition $Q \equiv \{x\}$, so property

$$Q \subseteq V_{ij}^i \text{ or } Q \subseteq V_{ij}^j \tag{5}$$

holds. Suppose it holds just before a subset A is tested on line 3 of Alg. 2. If $A \subseteq V_{ij}^i$, then by the induction assumption $\forall Q_i : Q_i \cap A \neq \emptyset, Q_i \subseteq V_{ij}^i$ and therefore $D \equiv \cup_{Q_i: Q_i \cap A \neq \emptyset} Q_i \subseteq V_{ij}^i$. Analogously with V_{ij}^j . This means that property (5) can only be violated after processing A if A intersects both V_{ij}^i and V_{ij}^j . In this case, from the definition of an ϵ -junction tree $I(V_{ij}^i, V_{ij}^j \mid S_{ij}) \leq$

ε , we have

$$\begin{aligned}
\varepsilon &\geq I(V_{ij}^i, V_{ij}^j | S_{ij}) \\
&\quad \{\text{monotonicity}\} \\
\varepsilon &\geq I(A \cap V_{ij}^i, A \cap V_{ij}^j | S_{ij}) \\
\varepsilon &\geq \min_{X \subset A} I(X, A-X | S_{ij}) \\
&\quad \{\delta = \varepsilon\} \\
\delta &\geq \min_{X \subset A} I(X, A-X | S_{ij})
\end{aligned}$$

which means that A is ignored by Alg 2 and property (5) holds throughout its runtime and in the end result.

Proof of $n(\varepsilon + (k-1)\delta)$ -weakness is also by induction. Using Lemma 4, it is enough to show that for every $A \subset V_{S_{ij}}$ s.t. $|A| \leq k+1$, $A \cap Q \neq \emptyset$ and $A \cap V_{S_{ij}Q} \neq \emptyset$ holds that

$$I(A \cap Q, A \cap V_{S_{ij}Q} | S_{ij}) \leq (k-1)\delta \quad (6)$$

Because Q remained disconnected from $V_{S_{ij}Q}$, and $|A| \leq k+1 \leq q$, we know that for every such A

$$\min_{X \subset A} I(X, A-X | S_{ij}) \leq \delta. \quad (7)$$

This means that for subsets A of size 2 equation (6) holds. Suppose it holds for A s.t. $|A| \leq m$. Let us prove that it has to hold for $m+1$ as well. Consider a subset A of size $m+1$. Denote

$$A^* \equiv \arg \min_{X \subset A} I(X, A-X | S_{ij})$$

Denote also (this notation is illustrated in Fig. 5)

$$W \equiv A^* \cap Q, \quad U \equiv A^* \cap V_{S_{ij}Q}, \quad Y \equiv A - A^* \cap Q, \quad Z \equiv A - A^* \cap V_{S_{ij}Q}.$$

From induction hypothesis, we know

$$I(W, U | S_{ij}) \leq (|A^*| - 1)\delta \text{ and } I(Y, Z | S_{ij}) \leq (|A| - |A^*| - 1)\delta = (m - |A^*|)\delta$$

and also from (7)

$$I(WU, YZ | S_{ij}) \leq \delta.$$

Let us show that

$$I(W, U | S_{ij}) + I(Y, Z | S_{ij}) + I(WU, YZ | S_{ij}) \geq I(WY, UZ | S_{ij})$$

Write down the difference in terms of entropies:

$$\begin{aligned}
&I(W, U | S_{ij}) + I(Y, Z | S_{ij}) + I(WU, YZ | S_{ij}) - I(WY, UZ | S_{ij}) \\
&= H(W | S_{ij}) - \mathbf{H}(\mathbf{W} | \mathbf{US}_{ij}) + H(Y | S_{ij}) - H(Y | ZS_{ij}) \\
&\quad + \mathbf{H}(\mathbf{WU} | \mathbf{S}_{ij}) - H(WU | YZS_{ij}) - H(WY | S_{ij}) + H(WY | UZS_{ij}) \\
&= \{\text{note that } H(WU | S_{ij}) - H(W | US_{ij}) = H(U | S_{ij})\} \\
&= H(W | S_{ij}) + H(U | S_{ij}) + H(Y | S_{ij}) - \mathbf{H}(\mathbf{Y} | \mathbf{ZS}_{ij}) \\
&\quad - \mathbf{H}(\mathbf{WU} | \mathbf{YZS}_{ij}) - H(WY | S_{ij}) + H(WY | UZS_{ij}) \\
&= \{\text{note that } H(WU | YZS_{ij}) + H(Y | ZS_{ij}) = H(WUY | ZS_{ij})\} \\
&= H(W | S_{ij}) + H(U | S_{ij}) + H(Y | S_{ij}) - \mathbf{H}(\mathbf{WUY} | \mathbf{ZS}_{ij}) - H(WY | S_{ij}) + \mathbf{H}(\mathbf{WY} | \mathbf{UZS}_{ij}) \\
&= \{\text{note that } H(WUY | ZS_{ij}) - H(WY | UZS_{ij}) = H(U | ZS_{ij})\} \\
&= H(W | S_{ij}) + H(U | S_{ij}) + H(Y | S_{ij}) - H(U | ZS_{ij}) - H(WY | S_{ij}) \\
&= I(W, Y | S_{ij}) + I(U, Z | S_{ij}) \geq 0
\end{aligned}$$

Therefore,

$$\begin{aligned}
I(WY, UZ | S_{ij}) &\leq I(W, U | S_{ij}) + I(Y, Z | S_{ij}) + I(WU, YZ | S_{ij}) \\
&\leq (|A^*| - 1)\delta + (m - |A^*|)\delta + \delta \\
&= m\delta = ((m+1) - 1)\delta.
\end{aligned}$$

But $I(WY, UZ \mid S_{ij}) \equiv I(A \cap Q, A \cap V_{-S_{ij}Q} \mid S_{ij})$ so (6) is proved. We can now apply Lemma 4, to show directly that

$$I(Q, V \setminus QS_{ij} \mid S_{ij}) \leq n(\varepsilon + (k-1)\delta).$$

□

Proof of Proposition 9. There are $\binom{n}{k}$ separators and at most $n-k$ components for each separator. For a component (S, Q) we need to consider $k|Q|$ possible subcomponents separators $S_i \in S \cup \{x\}$, $x \in Q$. The greedy checks of Def. 8 can be done in total time $O(n)$ for all components with a given separator S_i . The total complexity is therefore

$$O\left(\sum_{S:|S|=k} \sum_{Q:(S,Q) \in \mathbb{L}} k|Q|n\right) = O\left(\sum_{S:|S|=k} knn\right) = O\left(\binom{n}{k} kn^2\right) = O(kn^{k+2}) = O(n^{k+2})$$

(the last step is valid because k is considered to be a constant). □

Lemma 14. For non-intersecting sets A, B, S, Y, X and any probability distribution $P(ABSXY)$ it holds that

$$I(A, BY \mid SX) + I(AB, X \mid SY) \geq I(A, B \mid SY)$$

Proof. Write down the difference:

$$\begin{aligned} & I(A, BY \mid SX) + I(AB, X \mid SY) - I(A, B \mid SY) \\ &= H(ASX) - H(SX) - H(ABSXY) + H(BSXY) + \\ & \quad + H(ABSXY) - \mathbf{H}(\mathbf{SY}) - H(ABSXY) + H(SXY) - \\ & \quad - H(ASY) + \mathbf{H}(\mathbf{SY}) + H(ABSXY) - H(BSY) \\ &= \mathbf{H}(\mathbf{ASX}) - H(SX) - H(ABSXY) + H(BSXY) + H(ABSXY) \\ & \quad - H(ABSXY) + \mathbf{H}(\mathbf{SXY}) - H(ASY) + H(ABSXY) - H(BSY) \end{aligned}$$

We know that the entropy is submodular:

$$H(C) + H(D) \geq H(C \cap D) + H(C \cup D).$$

Let us use submodularity to bound the difference from below

$$\begin{aligned} & I(A, BY \mid SX) + I(AB, X \mid SY) - I(A, B \mid SY) \\ &\geq H(ASXY) + \mathbf{H}(\mathbf{SX}) - \mathbf{H}(\mathbf{SX}) - H(ABSXY) + H(BSXY) \\ & \quad + H(ABSXY) - H(ABSXY) - H(ASY) + H(ABSXY) - H(BSY) \\ &= \mathbf{H}(\mathbf{ASXY}) - H(ABSXY) + H(BSXY) \\ & \quad + \mathbf{H}(\mathbf{ABSXY}) - H(ABSXY) - H(ASY) + H(ABSXY) - H(BSY) \\ &\geq \mathbf{H}(\mathbf{ABSXY}) + \mathbf{H}(\mathbf{ASY}) - \mathbf{H}(\mathbf{ABSXY}) + H(BSXY) \\ & \quad - H(ABSXY) - \mathbf{H}(\mathbf{ASY}) + H(ABSXY) - H(BSY) \\ &= H(BSXY) - H(ABSXY) + H(ABSXY) - H(BSY) \\ &= I(A, SX \mid BSXY) \geq 0 \end{aligned}$$

□

Prof of Lemma 10. From the $n(\varepsilon + (k-1)\delta)$ -weakness property of Lemma 7, setting $\delta = \varepsilon$, we get that for every component of form $(S, Q) \in \mathbb{L}$ it holds that

$$I(Q, V_{-QS} \mid S) \leq nk\varepsilon. \quad (8)$$

Therefore, if Alg. 3 returns a junction tree, it will be a $nk\varepsilon$ -junction tree.

Let (T, \mathbb{C}) be the junction tree that satisfies the assumptions of Lemma 10. To show that Alg. 3 is guaranteed to return a tree, it is sufficient to show that for every component (S_{ij}, V_{ij}^i) from (T, \mathbb{C})

1. (S_{ij}, V_{ij}^i) will be in \mathbb{L} before *FindConsistentTreeDPGreedy* is invoked, and

2. *FindConsistentTreeDPGreedy* will find a decomposition for (S_{ij}, V_{ij}^i) .

Let us prove part 1. (T, \mathbb{C}) is an ε -junction tree, so for every (S_{ij}, V_{ij}^i) from (T, \mathbb{C}) it holds that

$$I(V_{ij}^i, V_{-S_{ij}V_{ij}^i} | S_{ij}) \leq \varepsilon,$$

and from correctness property of Lemma 7, we get that for every Q s.t. $(S_{ij}, Q) \in \mathbb{L}$ either $Q \subseteq V_{ij}^i$, or $Q \subseteq V_{ij}^j$. WLOG, assume that $Q \subseteq V_{ij}^i$. Suppose $Q \neq V_{ij}^i$. Denote $Q' \equiv V_{ij}^i \setminus Q$. Alg. 2 did not merge Q with Q' , so for every $A \subseteq Q, B \subseteq Q'$ s.t. $|AB| \leq k+2$ it holds that

$$\min_{X \subset AB} I(X, AB \setminus X | S_{ij}) \leq \varepsilon.$$

On the other hand, for every clique $C \in \mathbb{C}$ s.t. $C \subset V_{ij}^i \cup S_{ij}$ it holds that

$$\min_{X \subset (C \setminus S_{ij})} I(X, C \setminus X | S_{ij}) > (k+3)\varepsilon \geq \varepsilon, \quad (9)$$

so either $C \setminus S_{ij} \subseteq Q$ or $C \setminus S_{ij} \subseteq Q'$. Consider two cliques $C^1, C^2 \in \mathbb{C}$ such that there is an edge between C^1 and C^2 in (T, \mathbb{C}) and both C^1 and C^2 are on the same side of S_{ij} as variable i . The latter means in particular $C^1, C^2 \subset V_{ij}^i \cup S_{ij}$, and neither C^1 nor C^2 is C_j . We will show that either

$$C^1 \setminus S_{ij}, C^2 \setminus S_{ij} \subset Q$$

or

$$C^1 \setminus S_{ij}, C^2 \setminus S_{ij} \subset Q'.$$

Indeed, suppose neither is the case. Then

$$\begin{aligned} (C^1 \setminus S_{ij} \cap C^2 \setminus S_{ij}) &\subseteq (Q \cap Q') \equiv \emptyset \\ (C^1 \cap C^2) \setminus S_{ij} &= \emptyset. \\ S_{12} \setminus S_{ij} &= \emptyset. \end{aligned}$$

(T, \mathbb{C}) is maximal, so $|S_{12}| = |S_{ij}| = k$, and thus

$$S_{12} \equiv S_{ij}.$$

This is a contradiction with the condition of Lemma 10 that no two edges have the same separator, so either both $C^1 \setminus S_{ij}$ and $C^2 \setminus S_{ij}$ are in Q , or they are both in Q' .

Notice now that (T, \mathbb{C}) induces a tree over all cliques $C \in V_{ij}^i \cup S_{ij}$, so by transitive closure either $Q \equiv V_{ij}^i$ or $Q' \equiv V_{ij}^i$, so the assumption $Q \neq V_{ij}^i$ is incorrect. Therefore, all pairs of form (S_{ij}, V_{ij}^i) from (T, \mathbb{C}) are in \mathbb{L} before DP starts.

Now let us prove part 2 of the sufficient conditions. The proof is by induction on the size $|Q|$ of components (S, Q) . Consider a pair of separators S_{ij}, S_{im} from (T, \mathbb{C}) that are both connected to the clique C_i . Notice that

$$V_{ij}^i \supset V_{im}^m \Rightarrow |V_{ij}^i| > |V_{im}^m|.$$

Suppose a decomposition for (S_{im}, V_{im}^m) has been found. For leaf cliques C_m , $|V_{im}^m| = 1$ and the decomposition is trivial. The induction base is thus proven. When *FindConsistentTreeDPGreedy* checks for a decomposition of (S_{ij}, V_{ij}^i) , it will iterate through all variables x of V_{ij}^i , so at some point it will try $x = C_i \setminus S_{ij}$ and form a corresponding candidate clique $S_{ij} \cup x \equiv C_i$. Now *FindConsistentTreeDPGreedy* will iterate through the components that can be added to the decomposition of (S_{ij}, V_{ij}^i) . Component (S_{im}, V_{im}^m) satisfies the requirements $S_{im} \subseteq (S_{ij} \cup x) \equiv C_i$ and $V_{im}^m \subset V_{ij}^i$. Thus there are two possibilities. Either *FindConsistentTreeDPGreedy* adds (S_{im}, V_{im}^m) to the decomposition of (S_{ij}, V_{ij}^i) , or there is another component (S', Q) in the current (incomplete) decomposition of (S_{ij}, V_{ij}^i) such that it conflicts with (S_{im}, V_{im}^m) , i.e. $V_{im}^m \cap Q \neq \emptyset$. We will show that in the latter case $V_{im}^m \subseteq Q$.

Suppose $V_{im}^m \not\subseteq Q$, then in (T, \mathbb{C}) there exists a clique C on the same side of S_{im} as variable m s.t.

$$C \setminus S_{im} \cap Q \neq \emptyset, C \setminus S_{im} \setminus Q \neq \emptyset.$$

Notice that

$$(C_{-S_{im}} \cap Q) \cup (C_{-S_{im}} \setminus Q) \equiv C_{-S_{im}}$$

Denote

$$A \equiv C_{-S_{im}} \cap Q, B \equiv C_{-S_{im}} \setminus Q$$

By assumptions of the lemma, for a partitioning of $C \setminus S_{im}$ into sets A, B , it holds that $I(A, B \mid S_{im}) > (k + 3)\varepsilon$. Let us show that it follows that $I(A, B \mid S') > \varepsilon$. Denote

$$S_{ij} = (Sxy), S_{im} = (Syz), S' = (Sxz),$$

(then $C_i \equiv (Sxyz)$). From Lemma 14, we have

$$\begin{aligned} I(A, By \mid Sxz) &\geq I(A, B \mid Syz) - I(AB, x \mid Syz) \\ &> (k + 3)\varepsilon - I(AB, x \mid Syz) \\ &\quad \{Syz \text{ separates } x \text{ and } AB \text{ in } (T, \mathbb{C}), \text{ and } (T, \mathbb{C}) \text{ is an } \varepsilon\text{-JT}\} \\ &\geq (k + 3)\varepsilon - \varepsilon = (k + 2)\varepsilon \end{aligned} \tag{10}$$

On the other hand, we know that LTCI did not put A, B and y into the same partition Q for candidate separator Sxz . Because LTCI was called with max-set parameter $q = k + 2$, this means in particular that $\forall D \subseteq ABY$ it holds that

$$\min_{X \subseteq D} I(X, D_{-X} \mid Sxz) \leq \varepsilon.$$

By the same reasoning as in the proof of $n(\varepsilon + (k - 1)\delta)$ -weakness of LTCI in Lemma 7, it follows that

$$I(A, BY \mid Sxz) \leq (k + 2)\varepsilon,$$

a contradiction with Eqn. (10). Thus the assumption $V_{im}^m \not\subseteq Q$ is incorrect and in fact $V_{im}^m \subseteq Q$

Therefore, when *FindConsistentTreeDPGreedy* tries to add a true component (S_{im}, V_{im}^m) to the decomposition of (S_{ij}, V_{ij}^i) , either it succeeds, or all the variables from V_{im}^m are already covered by another component. $\cup_{m \neq j} V_{im}^m \equiv V_{ij}^i \setminus C_i$, so eventually all variables of $V_{ij}^i \setminus C_i$ will be covered, and a decomposition of (S_{ij}, V_{ij}^i) will thus be found. This holds for every separator from (T, \mathbb{C}) , so a junction tree will eventually be found \square

Proof of Theorem 12. The only part of the algorithm that is affected by the change of the oracle is Alg. 2. It is called $O(n^k)$ times, and each call performs $O(n^{k+2})$ oracle calls. Each oracle call is within $\pm\Delta$ of the true value with probability $1 - \gamma$. By union bound, with probability at least $1 - \gamma O(n^{2k+2})$ all of the oracle results will be within $\pm\Delta$ of their true values simultaneously.

For every (S, A) s.t. $I(A, V_{SA} \mid S) \leq \varepsilon$, it means that

$$\forall B : |B| \leq k + 2 \quad \hat{I}(B \cap A, B \cap V_{SA} \mid S) \leq \varepsilon + \Delta = \delta$$

and thus Alg. 3 will never introduce a link between A and V_{SA} . Thus correctness in the sense of Lemma 7 thus holds with $\delta = \varepsilon + \Delta$ for the result of Alg. 2.

Also, for any separator S , if $Q \subset V_S$ is in the set \mathbb{C}_S returned by Alg. 2, it holds that

$$\forall B : |B| \leq k + 2 \quad \hat{I}(B \cap Q, B \cap V_{SQ} \mid S) \leq \delta = \varepsilon + \Delta$$

therefore

$$I(B \cap Q, B \cap V_{SQ} \mid S) \leq \varepsilon + 2\Delta$$

so by the same reasoning as in the proof of Lemma 7, we have a modification of α -weakness property from that lemma.

$$I(A, V_{SA} \mid S) \leq n(\varepsilon + (k - 1)(\varepsilon + 2\Delta)) < kn(\varepsilon + 2\Delta)$$

Plugging the above results instead of Lemma 7 into the proof of Lemma 10, and adding Δ tolerances in appropriate places in that proof, we arrive to the statement of this theorem. \square

Proof of Corollary 13. Select $\Delta = \frac{\beta}{2(k+3)n}$. Since

$$2(k+3)\Delta = 2(k+3)\frac{\beta}{2(k+3)n} \leq 2(k+3)\frac{n\alpha}{2(k+3)n} = \alpha,$$

so an α -strongly connected junction tree is also $2(k+3)\Delta$ -connected, so from Thm. 12 Alg. 3 will learn a $2kn\Delta$ -junction tree using (domain size R is a constant)

$$O\left(\frac{n^2}{\Delta^2} \log^2 \frac{n}{\Delta} \log \frac{n}{\gamma}\right) = O\left(\frac{n^4}{\beta^2} \log^2 \frac{n}{\beta} \log \frac{n}{\gamma}\right)$$

samples and

$$O\left(\frac{n^{2k+5}}{\Delta^2} \log^2 \frac{n}{\Delta} \log \frac{n}{\gamma}\right) = O\left(\frac{n^{2k+7}}{\beta^2} \log^2 \frac{n}{\beta} \log \frac{n}{\gamma}\right)$$

computation time.

Since $2kn\Delta \leq \beta$, a $2kn\Delta$ -junction tree is also a β -junction tree. □