
On Theoretical Properties of Sum-Product Networks: Supplementary Material

Robert Peharz[†] **Sebastian Tschiatschek**[†]
[†]Signal Processing and
Speech Communication Lab,
Graz University of Technology

Franz Pernkopf[†] **Pedro Domingos**[‡]
[‡]Dept. of Computer Science & Engineering,
University of Washington

1 Examples

1.1 Example for partial evidence (Section 1)

Let $\mathbf{X} = \{X_1, X_2, X_3\}$ with $\mathbf{val}(X_1) = [0, 1]$, $\mathbf{val}(X_2) = \{1, 2, 3\}$ and $\mathbf{val}(X_3) = \mathbb{R}$. The partial evidence $\mathcal{X} = [0, 0.25] \cup [0.75, 1] \times \{1, 2, 3\} \times \{\pi\} \in \mathcal{H}_{\mathbf{X}}$ represents evidence that X_1 assumes a value smaller 0.25 or larger 0.75, X_2 assumes any of its possible states, and X_3 is set to the irrational number π .

1.2 Example for Algorithm 2

An example showing the mechanism of Algorithm 2 is shown in Figure 1. See the caption text for details.

1.3 Example for Algorithm 2 applied to a Sum-Product Tree

An example of applying Algorithm 2 to an SPT is shown in Figure 2, illustrating Proposition 3. See the caption text for details.

2 Proofs

Proposition 1. *Let \mathbf{P} be a product node and \mathbf{Y} be its shared RVs. \mathbf{P} is consistent iff for each $Y \in \mathbf{Y}$ there exists a unique $y^* \in \mathbf{val}(Y)$ with $\lambda_{Y=y^*} \in \mathbf{desc}(\mathbf{P})$.*

Proof. Direction “ \Leftarrow ”, i.e. suppose for each $Y \in \mathbf{Y}$ there exists a unique $y^* \in \mathbf{val}(Y)$ with $\lambda_{Y=y^*} \in \mathbf{desc}(\mathbf{P})$. Consider arbitrary two distinct children $C', C'' \in \mathbf{ch}(\mathbf{P})$ and let $\mathbf{Y}' = \mathbf{sc}(C') \cap \mathbf{sc}(C'')$, where from Definition 8 (shared RVs) it follows that $\mathbf{Y}' \subseteq \mathbf{Y}$. First, let $X \in \mathbf{sc}(C') \setminus \mathbf{Y}'$ arbitrary. Since $X \notin \mathbf{sc}(C'')$, we have

$$\lambda_{X=x} \in \mathbf{desc}(C') \Rightarrow \forall x' : \lambda_{X=x'} \notin \mathbf{desc}(C''), \quad (1)$$

and in particular

$$\lambda_{X=x} \in \mathbf{desc}(C') \Rightarrow \forall x' \neq x : \lambda_{X=x'} \notin \mathbf{desc}(C''). \quad (2)$$

Now let $X \in \mathbf{Y}'$ arbitrary. Since $X \in \mathbf{Y}$, there is a unique $y^* \in \mathbf{val}(Y)$ with $\lambda_{Y=y^*} \in \mathbf{desc}(C')$ and $\lambda_{Y=y^*} \in \mathbf{desc}(C'')$, and again (2) holds. Since C', C'' and X are arbitrary, \mathbf{P} is consistent.

Direction “ \Rightarrow ”, i.e. suppose \mathbf{P} is consistent. If \mathbf{Y} is empty, i.e. \mathbf{P} is decomposable, then the claimed property holds trivially for all \mathbf{Y} . Otherwise let $Y \in \mathbf{Y}$ arbitrary and let C' and C'' be arbitrary two children having Y in their scope. There are at least two such children by Definition 8 (shared RVs). There must be a unique $y^* \in \mathbf{val}(Y)$ with $\lambda_{Y=y^*} \in \mathbf{desc}(C')$ and $\lambda_{Y=y^*} \in \mathbf{desc}(C'')$. To see this, note that there must be at least one $y' \in \mathbf{val}(Y)$ with $\lambda_{Y=y'} \in \mathbf{desc}(C')$ and at least one $y'' \in \mathbf{val}(Y)$ with $\lambda_{Y=y''} \in \mathbf{desc}(C'')$, since Y is in the scope of both C' and C'' . However, there can be at most one such y' and y'' . Assume there were y'_1 and y'_2 , $y'_1 \neq y'_2$. Then either for $y' = y'_1$ or for $y' = y'_2$ we had a contradiction to

$$\lambda_{Y=y'} \in \mathbf{desc}(C') \Rightarrow \forall y'' \neq y' : \lambda_{Y=y''} \notin \mathbf{desc}(C''). \quad (3)$$

i.e. \mathbf{P} would not be consistent. By symmetry, there is also at most one y'' , and (3) can only hold when $y' = y'' =: y^*$. Therefore, since Y, C' and C'' are arbitrary, there must be a unique $y^* \in \mathbf{val}(Y)$ with $\lambda_{Y=y^*} \in \mathbf{desc}(C)$, $\forall C \in \mathbf{ch}(\mathbf{P}) : Y \in \mathbf{sc}(C)$. Therefore, this y^* is also the unique value with $\lambda_{Y=y^*} \in \mathbf{desc}(\mathbf{P})$. \square

Lemma 1. *Let \mathbf{N} be a node in some complete and consistent SPN over \mathbf{X} , $X \in \mathbf{sc}(\mathbf{N})$ and $x \in \mathbf{val}(X)$. When $\lambda_{X=x} \notin \mathbf{desc}(\mathbf{N})$, then $\forall \mathbf{x} \in \mathbf{val}(\mathbf{X})$ with $\mathbf{x}[X] = x$ we have $\mathbf{N}(\mathbf{x}) = 0$.*

Proof. When $\mathbf{N} = \lambda_{X=x'}$, for some $x' \in \mathbf{val}(X)$, the lemma clearly holds, since $x' \neq x$.

When \mathbf{N} is a product or a sum, let $\mathbf{N} = \{\mathbf{N}' \in \mathbf{desc}(\mathbf{N}) : X \in \mathbf{sc}(\mathbf{N}')\}$. Let $K = |\mathbf{N}|$ and let $\mathbf{N}_1, \dots, \mathbf{N}_K$ be a topologically ordered list of \mathbf{N} , i.e. $\mathbf{N}_k \notin \mathbf{desc}(\mathbf{N}_l)$ when $k > l$. We can assume that for some I , $\mathbf{N}_1, \dots, \mathbf{N}_I$ are IVs of X , $\mathbf{N}_{I+1}, \dots, \mathbf{N}_K$ are sum and product nodes, where

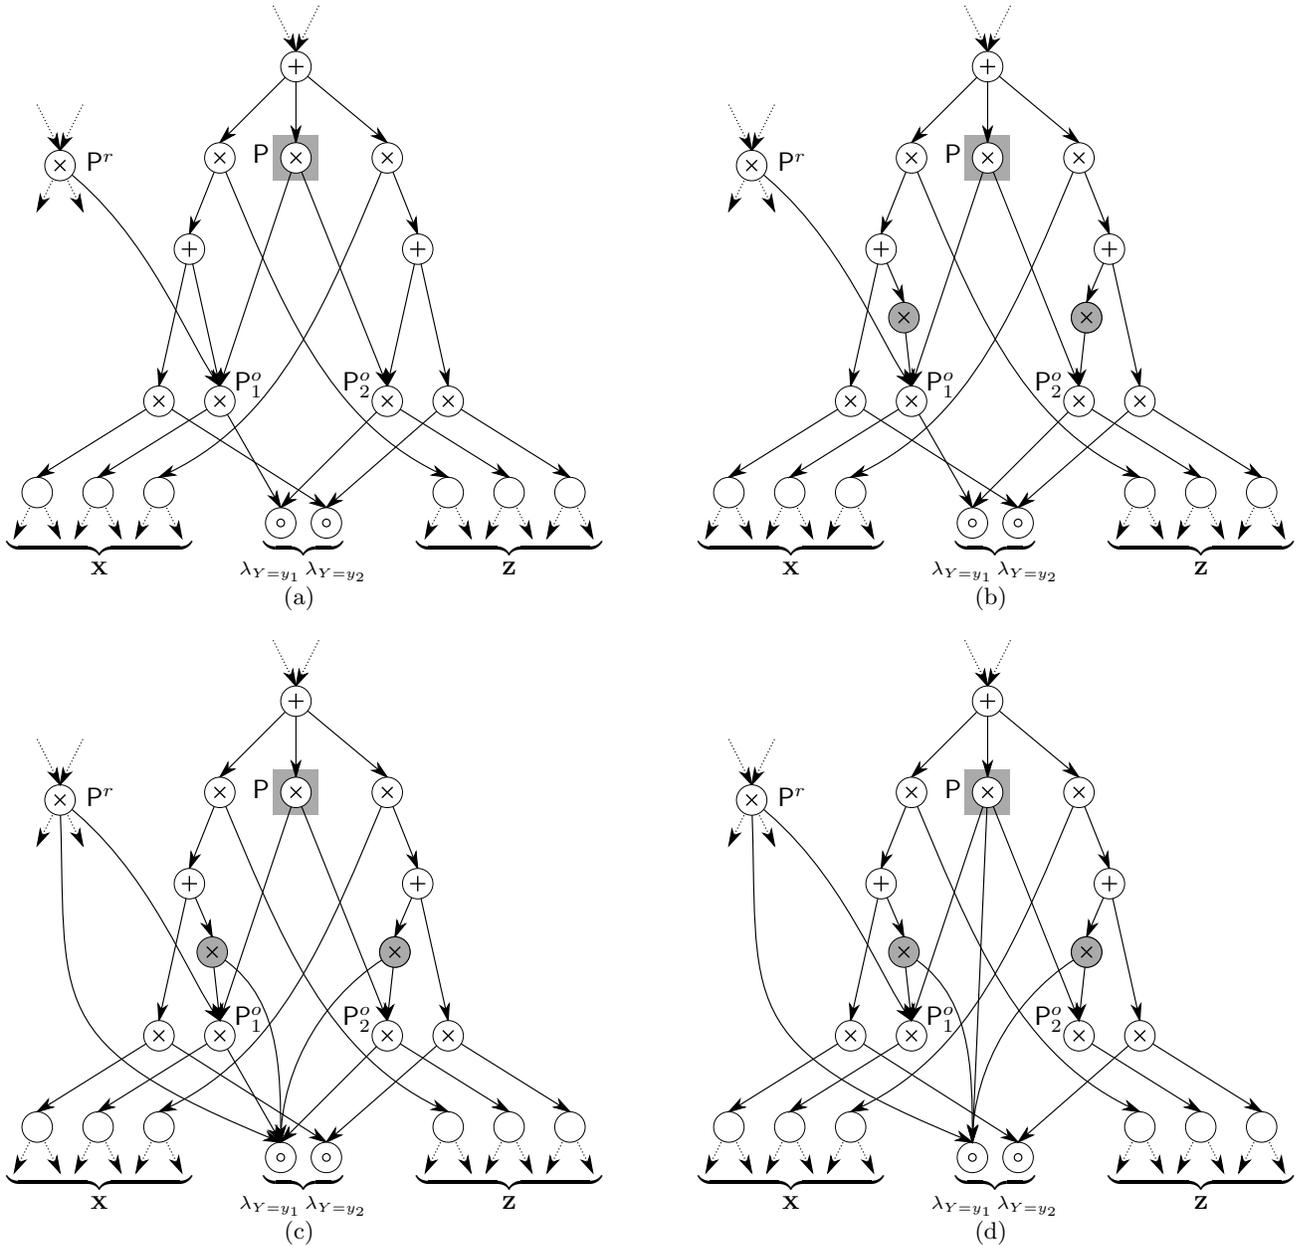


Figure 1: Illustration of Algorithm 2, transforming a complete and consistent SPN into a complete and decomposable one. (a): Excerpt of an SPN containing a single consistent, non-decomposable product P with shared RVs $Y = \{Y\}$, since $\lambda_{Y=y_1}$ is reached via both P_1^o and P_2^o . Dotted edges denote a continuation of the SPN outside this excerpt. The blank circles at the bottom symbolize sub-SPNs over RV sets X and Z , where we assume that $X \cap Z = \emptyset$. P 's children P_1^o and P_2^o have both a sum node as co-parent. P_1^o has additionally a product P^r as co-parent in some remote part of the SPN. (b): Introducing links, depicted as grey product nodes, according to steps 2–7 of Algorithm 2. (c): Steps 15–20. Here, the set \mathbf{N}^o is given as $\mathbf{N}^o = \{P, P_1^o, P_2^o\}$. All parents of $\mathbf{N}^o \setminus \{P\}$ which are not descendants of P , i.e. the two links and P^r , are connected with the IVs of the respective part of the shared RVs Y , here always $\lambda_{Y=y_1}$. (d): Rendering P decomposable, steps 21–24. IV $\lambda_{Y=y_1}$ is disconnected from P_1^o, P_2^o in steps 21–23, cutting Y from P_1^o, P_2^o and P . Step 24: IV $\lambda_{Y=y_1}$ is re-connected to P , which computes the same distribution as before, but is decomposable now. P_1^o and P_2^o could be short-wired and removed, since they are left with only one child. However, this does not necessarily happen and does not improve the theoretical bound of additionally required multiplications.

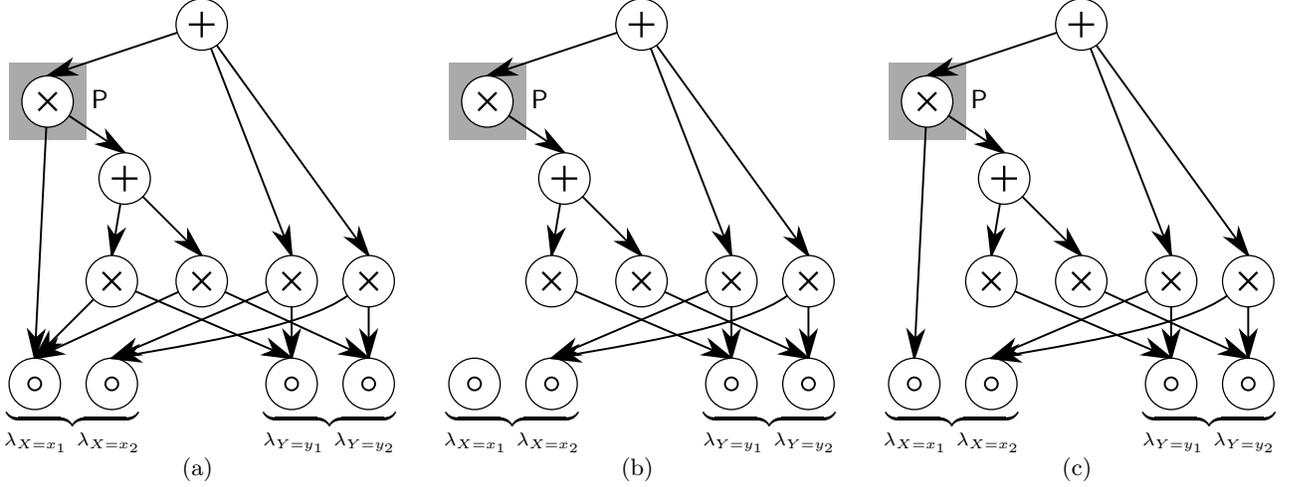


Figure 2: Illustration of Algorithm 2, transforming a complete and consistent, but not decomposable SPT into a complete and decomposable one. This demonstrates that using non-decomposable products is wasteful in SPTs (cf. Proposition 3), since after performing Algorithm 2, the SPT computes the same distribution, but requires fewer multiplications. (a): Complete and consistent SPT over two binary RVs X, Y , containing a non-decomposable product P . (b): Effect of steps 21–23 of Algorithm 2, saving 3 multiplications. (c): Effect of steps 24 of Algorithm 2, reconnecting $\lambda_{X=x_1}$ to P , re-introducing 1 multiplication.

$N_K = N$. For any $\mathbf{x} \in \text{val}(\mathbf{X})$ with $\mathbf{x}[X] = x$, we have $N_1(\mathbf{x}) = N_2(\mathbf{x}) = \dots = N_I(\mathbf{x}) = 0$, since $\lambda_{X=x} \notin \text{desc}(\mathbf{N})$.

When all $N_1(\mathbf{x}) = N_2(\mathbf{x}) = \dots = N_k(\mathbf{x}) = 0$ for some $k \geq I$, then also $N_{k+1}(\mathbf{x}) = 0$. When N_{k+1} is a sum node, due to completeness, all children of N_{k+1} must have X in their scope. This means that all children are in $\{N_1, \dots, N_k\}$. Therefore $N_{k+1}(\mathbf{x}) = 0$. When N_{k+1} is a product, at least one child has to be in $\{N_1, \dots, N_k\}$. Therefore $N_{k+1}(\mathbf{x}) = 0$. Thus, by induction, for all $N' \in \mathbf{N}$ we have $N'(\mathbf{x}) = 0$. In particular this is true for $N_K = N$. \square

Lemma 2. Let P be a probability mass function (PMF) over \mathbf{X} and $\mathbf{Y} \subseteq \mathbf{X}$, $\mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$ such that there exists a $\mathbf{y}^* \in \text{val}(\mathbf{Y})$ with $P(\mathbf{z}, \mathbf{y}) = 0$ when $\mathbf{y} \neq \mathbf{y}^*$. Then we have $P(\mathbf{z}, \mathbf{y}) = \mathbb{1}(\mathbf{y} = \mathbf{y}^*) P(\mathbf{z})$.

Proof. Since $P(\mathbf{z}, \mathbf{y}) = 0$ for $\mathbf{y} \neq \mathbf{y}^*$, $P(\mathbf{z}) = \sum_{\mathbf{y} \in \text{val}(\mathbf{Y})} P(\mathbf{z}, \mathbf{y}) = P(\mathbf{z}, \mathbf{y}^*)$. Thus

$$P(\mathbf{z}, \mathbf{y}) = \begin{cases} 0 & \text{if } \mathbf{y} \neq \mathbf{y}^* \\ P(\mathbf{z}) & \text{if } \mathbf{y} = \mathbf{y}^* \end{cases} = \mathbb{1}(\mathbf{y} = \mathbf{y}^*) P(\mathbf{z}). \quad (4)$$

\square

Theorem 1. Let \mathcal{S} be a complete and consistent SPN and P be a non-decomposable product in \mathcal{S} , \mathbf{Y} be the shared RVs of P and \mathbf{y}^* the consistent state of \mathbf{Y} . For $N \in \text{desc}(P)$ define $\mathbf{Y}^N := \mathbf{Y} \cap \text{sc}(N)$ and $\mathbf{X}^N := \text{sc}(N) \setminus \mathbf{Y}^N$. Then for all $N \in \text{desc}(P)$, and all $\mathbf{x} \in$

$\text{val}(\text{sc}(N))$:

$$P_N(\mathbf{x}) = \mathbb{1}(\mathbf{x}[\mathbf{Y}^N] = \mathbf{y}^*[\mathbf{Y}^N]) P_N(\mathbf{x}[\mathbf{X}^N]), \quad (5)$$

where $\mathbb{1}$ is the indicator function.

Proof. From Proposition 1 we know that $\forall Y \in \mathbf{Y}: \lambda_{Y=\mathbf{y}^*[Y]} \in \text{desc}(P)$ and $\forall y \neq \mathbf{y}^*[Y]: \lambda_{Y=y} \notin \text{desc}(P)$. Consequently, for any $N \in \text{desc}(P)$ we have for all $Y \in \mathbf{Y}^N$ that $\lambda_{Y=\mathbf{y}^*[Y]} \in \text{desc}(N)$ and $\forall y \neq \mathbf{y}^*[Y]: \lambda_{Y=y} \notin \text{desc}(N)$. With Lemma 1 it follows that for all $\mathbf{x} \in \text{val}(\text{sc}(N))$ with $\mathbf{x}[\mathbf{Y}^N] \neq \mathbf{y}^*[\mathbf{Y}^N]$, we have $P_N(\mathbf{x}) = 0$. Theorem 1 follows with Lemma 2. \square

Corollary 1. Let $P, \mathbf{Y}, \mathbf{y}^*$ be as in Theorem 1. For $C \in \text{ch}(P)$, let $\mathbf{X}^C := \text{sc}(C) \setminus \mathbf{Y}$, i.e. the part of $\text{sc}(C)$ which belongs exclusively to C . Then

$$P_P(\mathbf{x}) = \mathbb{1}(\mathbf{x}[\mathbf{Y}] = \mathbf{y}^*) \prod_{C \in \text{ch}(P)} C(\mathbf{x}[\mathbf{X}^C]). \quad (6)$$

Theorem 2. For each complete and consistent SPN $\mathcal{S}' = (\mathcal{G}', \mathbf{w}')$, there exists a complete, consistent and locally normalized SPN $\mathcal{S} = (\mathcal{G}, \mathbf{w})$ with $\mathcal{G}' = \mathcal{G}$, such that $\forall N \in \mathcal{G}: \mathcal{S}_N \equiv P_{\mathcal{S}'_N}$.

Proof. Algorithm 1 finds locally normalized weights without changing the distribution of any node. For deriving the algorithm, we introduce a correction factor α_P for each product node P , initialized to $\alpha_P = 1$. We redefine the product node P as $P(\lambda) := \alpha_P P(\lambda)$. At the end of the algorithm, all α_P will be 1 again.

Algorithm 1 Locally Normalize SPN

```

1: Let  $N_1, \dots, N_K$  be a topologically ordered list of
   all sum and product nodes
2: For all product nodes  $P$  initialize  $\alpha_P \leftarrow 1$ 
3: for  $k = 1 : K$  do
4:   if  $N_k$  is a sum node then
5:      $\alpha \leftarrow \sum_{C \in \text{ch}(N_k)} w_{N_k, C}$ 
6:      $\forall C \in \text{ch}(N_k) : w_{N_k, C} \leftarrow \frac{w_{N_k, C}}{\alpha}$ 
7:   end if
8:   if  $N_k$  is a product node then
9:      $\alpha \leftarrow \alpha_{N_k}$ 
10:     $\alpha_{N_k} \leftarrow 1$ 
11:   end if
12:   for  $F \in \text{pa}(N_k)$  do
13:     if  $F$  is a sum node then
14:        $w_{F, N_k} \leftarrow \alpha w_{F, N_k}$ 
15:     end if
16:     if  $F$  is a product node then
17:        $\alpha_F \leftarrow \alpha \alpha_F$ 
18:     end if
19:   end for
20: end for
    
```

Let N'_1, \dots, N'_K be a topologically ordered list of all nodes in the unnormalized S' , i.e. $N'_k \notin \text{desc}(N'_l)$ if $k > l$. Let N_1, \dots, N_K be the corresponding list of \mathcal{S} , which will be the normalized version after the algorithm has terminated. We have the following loop invariant for the main loop. Given that at the k^{th} entrance of the main loop

1. $P_{N'_l} = \mathcal{S}_{N_l}$, for $1 \leq l < k$
2. $S'_{N'_m} = \mathcal{S}_{N_m}$, for $k \leq m \leq K$

the same will hold for $k + 1$ at the end of the loop.

The first point holds since we normalize N_k during the main loop: All nodes prior in the topological order, and therefore all children of N_k are already normalized. If N_k is a sum, then it represents a mixture distribution after step 6. If N_k is a product, then it will be normalized after step 10 since we set $\alpha_{N_k} = 1$.

The second point holds since the modification of N_k can change any N_m , $m > k$, only via $\text{pa}(N_k)$. The change of N_k is compensated for all its parents either in step 14 or step 17, depending on whether the parent is a sum or a product node.

From this loop invariance it follows by induction that all N_1, \dots, N_K compute the normalized distributions of N'_1, \dots, N'_K after the K^{th} iteration. \square

Proposition 2. *A complete and decomposable SPN computes the NP of some unnormalized distribution.*

Proof. Assume some topological ordering of the SPN nodes. We show by induction over this order, from bottom to top, that each sub-SPN \mathcal{S}_N computes an NP over $\text{sc}(N)$. The induction basis are the IVs $\lambda_{X=x}$ which compute

$$\lambda_{X=x} = \sum_{x' \in \text{val}(X)} \mathbf{1}(x' = x) \lambda_{X=x'}, \quad (7)$$

which is an NP over X .

A complete sum node S computes

$$S(\lambda) = \sum_{C \in \text{ch}(S)} w_{S, C} C(\lambda) \quad (8)$$

$$= \sum_{C \in \text{ch}(S)} w_{S, C} \sum_{\mathbf{x} \in \text{val}(\text{sc}(S))} \Phi_C(\mathbf{x}) \prod_{X \in \text{sc}(S)} \lambda_{X=\mathbf{x}[X]} \quad (9)$$

$$= \sum_{\mathbf{x} \in \text{val}(\text{sc}(S))} \Phi_S(\mathbf{x}) \prod_{X \in \text{sc}(S)} \lambda_{X=\mathbf{x}[X]}, \quad (10)$$

where Φ_C is the unnormalized distributions of the NP of C , and $\Phi_S(\mathbf{x}) := \sum_{C \in \text{ch}(S)} w_{S, C} \Phi_C(\mathbf{x})$. In (9) we apply the induction hypothesis and the fact that for complete sum nodes the scopes of the node and all its children are the same. We see that (10) has the form of an NP, i.e. the induction step holds for complete sum nodes.

Now consider some decomposable product node P with children $\text{ch}(P) = \{C_1, \dots, C_K\}$. The product node P computes

$$P(\lambda) = \prod_{k=1}^K C_k(\lambda) \quad (11)$$

$$= \prod_{k=1}^K \sum_{\mathbf{x}^k \in \text{val}(\text{sc}(C_k))} \Phi_{C_k}(\mathbf{x}^k) \prod_{X \in \text{sc}(C_k)} \lambda_{X=\mathbf{x}^k[X]} \quad (12)$$

$$= \sum_{\mathbf{x}^1} \dots \sum_{\mathbf{x}^K} \prod_{k=1}^K \left[\Phi_{C_k}(\mathbf{x}^k) \prod_{X \in \text{sc}(C_k)} \lambda_{X=\mathbf{x}^k[X]} \right] \quad (13)$$

$$= \sum_{\mathbf{x} \in \text{val}(\text{sc}(P))} \Phi_P(\mathbf{x}) \prod_{X \in \text{sc}(P)} \lambda_{X=\mathbf{x}[X]}, \quad (14)$$

where Φ_{C_k} is the distributions of the NP of C_k , and $\Phi_P(\mathbf{x}) := \prod_{k=1}^K \Phi_{C_k}(\mathbf{x}[\text{sc}(C_k)])$. In (12) we use the induction hypothesis, in (13) we apply the distributive law, and in (14) we use the fact that the scope of a decomposable product node is partitioned by the scopes of its children. We see that (14) has the form of an NP, i.e. the induction step holds for complete product nodes.

Consequently, each sub-SPN and therefore also the overall SPN computes an NP of some unnormalized distribution. \square

Theorem 3. *Every complete and consistent SPN $\mathcal{S} = ((V, E), \mathbf{w})$ over \mathbf{X} can be transformed into a complete and decomposable SPN $\mathcal{S}' = ((V', E'), \mathbf{w}')$ over \mathbf{X} such that $P_{\mathcal{S}} \equiv P_{\mathcal{S}'}$, and where $|V'| \in \mathcal{O}(|V|^2)$, $A_{\mathcal{S}'} = A_{\mathcal{S}}$ and $M_{\mathcal{S}'} \in \mathcal{O}(M_{\mathcal{S}} |\mathbf{X}|)$.*

Algorithm 2 Transform to decomposable SPN

```

1: Let  $\mathbf{N} = N_1, \dots, N_K$  be a topologically ordered list
   of all sums and products
2: for all sum nodes  $S$  and all  $C \in \text{ch}(S)$  do
3:   if  $\text{pa}(C) > 1$  then
4:     Generate a new product node  $P^{S,C}$ 
5:     Interconnect  $P^{S,C}$  between  $S$  and  $C$ 
6:   end if
7: end for
8: while exist non-decomposable products in  $\mathbf{N}$  do
9:    $P \leftarrow N_{\min\{k' \mid N_{k'} \text{ is a non-decomposable product}\}}$ 
10:   $\mathbf{Y} \leftarrow$  shared RVs of  $P$ 
11:   $\mathbf{y}^* \leftarrow$  consistent state of  $\mathbf{Y}$ 
12:  if  $\text{sc}(P) = \mathbf{Y}$  then
13:    Replace  $P$  by  $\prod_{Y \in \mathbf{Y}} \lambda_{Y=\mathbf{y}^*}[Y]$ 
14:  else
15:     $\mathbf{N}^d \leftarrow$  sums and products in  $\text{desc}(P)$ 
16:     $\mathbf{N}^o \leftarrow \{N \in \mathbf{N}^d : \text{sc}(N) \not\subseteq \mathbf{Y}, \text{sc}(N) \cap \mathbf{Y} \neq \emptyset\}$ 
17:    for  $N \in \mathbf{N}^o \setminus \{P\}$  do
18:       $\mathbf{F} \leftarrow \text{pa}(N) \setminus \mathbf{N}^d$ 
19:       $\forall Y \in \mathbf{Y} \cap \text{sc}(N)$ :
20:        connect  $\lambda_{Y=\mathbf{y}^*}[Y]$  as child of all  $\mathbf{F}$ 
21:    end for
22:    for  $P^o \in \mathbf{N}^o$  do
23:      Disconnect  $C \in \text{ch}(P^o)$  if  $\text{sc}(C) \subseteq \mathbf{Y}$ 
24:    end for
25:     $\forall Y \in \mathbf{Y}$ : connect  $\lambda_{Y=\mathbf{y}^*}[Y]$  as child of  $P$ 
26:  end while
27: Delete all unreachable sums and products

```

Proof. Due to Theorem 2 we assume w.l.o.g. that \mathcal{S}' is locally normalized, and thus $P_{\mathcal{S}} \equiv \mathcal{S}$. Algorithm 2 transforms \mathcal{S} into a complete and decomposable SPN, representing the same distribution. First it finds a topologically ordered list N_1, \dots, N_K of all sum and product nodes, i.e. $k > l \Rightarrow N_k \notin \text{desc}(N_l)$. Then, in steps 2–7, it considers all sum nodes S and all children $C \in \text{ch}(S)$; if the child C has further parents except S , a newly generated product node $P^{S,C}$ is interconnected between S and C , i.e. $P^{S,C}$ is connected as child of S with weight $w_{S,C}$, C is disconnected from S and connected as child of $P^{S,C}$. To $P^{S,C}$ we refer as *link* between S and C . Note that the link has *only* S as parent, i.e. the link represents a *private* copy of child C for sum node S . Clearly, after step 7, the SPN still computes the same function.

In each iteration of the main loop 8–26, the algorithm

finds the lowest non-decomposable product node $N_k = P$ w.r.t. the topological ordering. We distinguish two cases: $\text{sc}(P) = \mathbf{Y}$ and $\text{sc}(P) \neq \mathbf{Y} \Leftrightarrow \mathbf{Y} \subset \text{sc}(P)$.

In the first case, we know from Corollary 1 that $P(\mathbf{y}) = \mathbb{1}(\mathbf{y} = \mathbf{y}^*)$, which is equivalent to the decomposable product $\prod_{Y \in \mathbf{Y}} \lambda_{Y=\mathbf{y}^*}[Y]$ replacing P , i.e. this new product is connected as child of all parents of P , and P itself is deleted. Deletion of P might render some nodes unreachable; however, these unreachable nodes do not “influence” the root node and will be safely deleted in step 27.

In the second case, when $\mathbf{Y} \subset \text{sc}(P)$, the algorithm first finds the set \mathbf{N}^d of all sum and product descendants of P . It also finds the subset \mathbf{N}^o of \mathbf{N}^d , containing all nodes whose scope overlaps with \mathbf{Y} , but is no subset of \mathbf{Y} . Clearly, P is contained in \mathbf{N}^o . The basic strategy is to “cut” \mathbf{Y} from the scope of P , i.e. that \mathbf{Y} is marginalized, rendering P decomposable. Then, by re-connecting all indicators $\lambda_{Y=\mathbf{y}^*}[Y]$ to P in step 24, P computes the same distribution as before due to Corollary 1, but is rendered *decomposable* now. Steps 21–23 cut \mathbf{Y} from *all* nodes in \mathbf{N}^o , in particular from P , but leave all sub-SPNs rooted at any node in $\mathbf{N}^d \setminus \mathbf{N}^o$ unchanged. To see this, note that $\mathbf{N}^d \setminus \mathbf{N}^o$ contains two types of nodes:

- nodes \mathbf{N}^s whose scope is a subset of \mathbf{Y} , i.e. $\mathbf{N}^s = \{N \in \mathbf{N}^d \mid \text{sc}(N) \subseteq \mathbf{Y}\}$, and
- nodes \mathbf{N}^n whose scope does not contain any \mathbf{Y} , i.e. $\mathbf{N}^n = \{N \in \mathbf{N}^d \mid \text{sc}(N) \cap \mathbf{Y} = \emptyset\}$.

Clearly, sub-SPNs rooted at any node in \mathbf{N}^s or \mathbf{N}^n do not contain any \mathbf{N}^o . Steps 21–23 only delete some outgoing edges of some $P^o \in \mathbf{N}^o$; thus all sub-SPNs rooted at nodes in \mathbf{N}^s or \mathbf{N}^n remain unchanged, and do not change the output of the overall SPN via nodes outside of \mathbf{N}^d . Now consider a topologically ordered list N_1^o, \dots, N_L^o of the nodes in \mathbf{N}^o , i.e. $k > l \Rightarrow N_k^o \notin \text{desc}(N_l^o)$. N_1^o must be a product. If it was a sum, all its children would have the same scope as N_1^o , due to completeness. Therefore, all its children would be contained in \mathbf{N}^o , contradicting that N_1^o is the first node in the topological order. Thus N_1^o is a product. N_1^o can have three types of children: nodes in \mathbf{N}^s , nodes in \mathbf{N}^n and IVs. Nodes in \mathbf{N}^s and IVs corresponding to some $Y \in \mathbf{Y}$ are disconnected from product N_1^o in steps 21–23. These children make up the deterministic term $\mathbb{1}(\cdot)$ in Theorem 1. By disconnecting them, $\text{sc}(N_1^o) \cap \mathbf{Y}$ is “cut” from N_1^o , which now computes the marginal distribution over $\text{sc}(N_1^o) \setminus \mathbf{Y}$. This is the induction basis. Now assume that after steps 21–23 all N_1^o, \dots, N_l^o , $l < L$, compute the marginal, \mathbf{Y} being marginalized. Then also N_{l+1}^o computes such a marginal. If N_{l+1}^o is a sum and thus a mixture distribution, this clearly

holds, since mixture sums and marginalization sums can be swapped, i.e. “the mixture of marginals is the marginal of the mixture”. If \mathbf{N}_{l+1}^o is a product, it can have four types of children: nodes in \mathbf{N}^o , nodes in \mathbf{N}^s , nodes in \mathbf{N}^n and IVs. By induction hypothesis \mathbf{Y} is already cut from all children in \mathbf{N}^o . Nodes in \mathbf{N}^s and IVs corresponding to some $Y \in \mathbf{Y}$ are disconnected from product \mathbf{N}_l^o . These nodes make up the deterministic term $\mathbb{1}(\cdot)$ in Theorem 1; a subset of \mathbf{Y} might already have been removed, since \mathbf{Y} has been removed from all children in \mathbf{N}^o . By disconnecting children in \mathbf{N}^s and IVs corresponding to \mathbf{Y} , this deterministic term is fully removed, and \mathbf{N}_l^o now computes the marginal distribution over $\mathbf{sc}(\mathbf{N}_l^o) \setminus \mathbf{Y}$. Thus, by induction, after steps 21–23 all nodes in \mathbf{N}^o compute the marginal distribution, \mathbf{Y} being marginalized.

Although we achieve our primary goal to render \mathbf{P} decomposable, steps 21–23 also cut \mathbf{Y} from any other node in $\mathbf{N} \in \mathbf{N}^o$, which would modify the SPN output via \mathbf{N} ’s parents *outside* of \mathbf{N}^d , i.e. via $\mathbf{F} = \mathbf{pa}(\mathbf{N}) \setminus \mathbf{N}^d$. Note that all nodes in \mathbf{F} must be products. To see this, assume that \mathbf{F} contains a sum \mathbf{S} . This would imply that \mathbf{N} is a link, which can be reached from \mathbf{P} only via its single parent \mathbf{S} . This implies $\mathbf{S} \in \mathbf{N}^d$, a contradiction. By Theorem 1, the distribution of \mathbf{N} is deterministic w.r.t. $\mathbf{sc}(\mathbf{N}) \cap \mathbf{Y}$. Steps 21–23 cut \mathbf{Y} from \mathbf{N} , which would change the distribution of the nodes in \mathbf{F} . Thus, in step 19 the IVs corresponding to $\mathbf{Y} \cap \mathbf{sc}(\mathbf{N})$ are connected to all \mathbf{F} , such that they still “see” the same distribution after steps 21–23. It is easy to see that if some $\mathbf{F} \in \mathbf{F}$ was decomposable beforehand, it will also be decomposable after step 23, i.e. steps 15–24 do not render other products non-decomposable.

Since the only new introduced nodes are the links between sum nodes and their children, and the number of sum-edges is in $\mathcal{O}(|V|^2)$, we have $|V'| \in \mathcal{O}(|V|^2)$. The number of summations is the same in \mathcal{S} and \mathcal{S}' , i.e. $A_{\mathcal{S}'} = A_{\mathcal{S}}$. Furthermore, introducing the links can not introduce more than double the number of multiplications, since we already require one multiplication per sum-edge. Thus, after step 7 we have $M_{\mathcal{S}'} \in \mathcal{O}(M_{\mathcal{S}})$. Since the while-loop in Algorithm 2 can not connect more than one IV per $X \in \mathbf{X}$ to each product node, we have $M_{\mathcal{S}'} \in \mathcal{O}(M_{\mathcal{S}} |\mathbf{X}|)$. \square

Proposition 3. *Every complete and consistent, but non-decomposable SPT $\mathcal{S} = ((V, E), \mathbf{w})$ over \mathbf{X} can be transformed into a complete and decomposable SPT $\mathcal{S}' = ((V', E'), \mathbf{w}')$ over \mathbf{X} such that $P_{\mathcal{S}} \equiv P_{\mathcal{S}'}$, and where $|V'| \leq |V|$, $A_{\mathcal{S}'} \leq A_{\mathcal{S}}$ and $M_{\mathcal{S}'} < M_{\mathcal{S}}$.*

Proof. For SPTs, no links are introduced in steps 2–7 of Algorithm 2. Consider a non-decomposable product \mathbf{P} in the SPT, and let \mathbf{Y} be the shared RVs. If $\mathbf{sc}(\mathbf{P}) = \mathbf{Y}$, the whole sub-SPN rooted at \mathbf{P} will be replaced

by $\prod_{Y \in \mathbf{Y}} \lambda_{Y=y^*[Y]}$ in step 13. The deleted sub-SPN computes $\prod_{Y \in \mathbf{Y}} (\lambda_{Y=y^*[Y]})^{k_Y}$, where k_Y are integers with $k_Y \geq 1$, and at least one $k_Y > 1$, since otherwise \mathbf{P} would be decomposable. Thus, replacing this sub-SPN by the decomposable product in step 13 saves at least one multiplication.

Similarly, when $\mathbf{sc}(\mathbf{P}) \neq \mathbf{Y}$, steps 21–23 prune nodes from the sub-SPN rooted at \mathbf{P} , corresponding to the computation of $\prod_{Y \in \mathbf{Y}} (\lambda_{Y=y^*[Y]})^{k_Y}$ with $k_Y \geq 1$, and at least one $k_Y > 1$. This requires at least one multiplication more than directly connecting $\lambda_{Y=y^*[Y]}$ to \mathbf{P} in step 24. Clearly, steps 17–20 are never performed in SPTs, since no $\mathbf{N} \in \mathbf{desc}(\mathbf{P}) \setminus \{\mathbf{P}\}$ can have a parent outside the descendants of \mathbf{P} .

Thus, for every non-decomposable product, Algorithm 2 saves at least one multiplication. \square

Proposition 4. *Let \mathcal{S}^g be a gated SPN. For any $X \in \mathbf{X}$ and $k \in \{1, \dots, |\mathbf{D}^X|\}$, let $\mathbf{X}^D = \mathbf{sc}(\mathbf{D}^{X,k})$ and $\mathbf{X}^R = \mathbf{X} \setminus \mathbf{X}^D$. It holds that $\mathcal{S}^g(\mathbf{X}, Z_X = k, \mathbf{Z} \setminus Z_X) = \mathbf{D}^{X,k}(\mathbf{X}^D) \mathcal{S}^g(\mathbf{X}^R, Z_X = k, \mathbf{Z} \setminus Z_X)$.*

Proof. Consider all sum and product nodes \mathbf{N} in \mathcal{S}^g which have X in their scope. We can split \mathbf{N} into two sets \mathbf{N}' , \mathbf{N}'' , where

$$\begin{aligned} \mathbf{N}' &= \{\mathbf{N} \in \mathbf{N} \mid \mathbf{D}^{X,k} \in \mathbf{desc}(\mathbf{N})\} \\ &= \{\mathbf{N} \in \mathbf{N} \mid \lambda_{Z=X,k} \in \mathbf{desc}(\mathbf{N})\} \end{aligned} \quad (15)$$

and $\mathbf{N}'' = \mathbf{N} \setminus \mathbf{N}'$. When we set $Z_X = k$, we get by Lemma 1 that all $\mathbf{N} \in \mathbf{N}''$ output 0.

We now show by induction that all $\mathbf{N} \in \mathbf{N}'$ have the form

$$\mathbf{D}^{X,k}(\mathbf{X}^D) \mathbf{N}(\mathbf{X}^{R'}, Z_X = k, \mathbf{Z}' \setminus \{Z_X\}), \quad (16)$$

where $\mathbf{X}^{R'} = \mathbf{X}^R \cap \mathbf{sc}(\mathbf{N})$ and $\mathbf{Z}' = \mathbf{Z} \cap \mathbf{sc}(\mathbf{N})$. Let $\mathbf{N}_1, \dots, \mathbf{N}_K$ be a topologically ordered list of \mathbf{N}' , i.e. $\mathbf{N}_k \notin \mathbf{desc}(\mathbf{N}_l)$ when $k > l$, where \mathbf{N}_1 is the introduced product when constructing the gated SPN, having $\mathbf{D}^{X,k}$ and $\lambda_{Z_X=k}$ as children, and \mathbf{N}_K is the root node. Eq. (16) clearly holds for \mathbf{N}_1 , since $\mathbf{N}_1(\mathbf{X}^D, Z_X = k) = \mathbf{D}^{X,k}(\mathbf{X}^D) \lambda_{Z_X=k}(Z_X = k)$. This is the induction basis. For the induction step assume that (16) holds for $\mathbf{N}_1, \dots, \mathbf{N}_k$. When \mathbf{N}_{k+1} is a sum

S, it computes

$$\sum_{C \in \text{ch}(S)} w_{S,C} C(\mathbf{X}, Z_X = k, \mathbf{Z}' \setminus \{Z_X\}) \quad (17)$$

$$\begin{aligned} &= \sum_{C \in \text{ch}(S) \cap \mathbf{N}'} w_{S,C} C(\mathbf{X}, Z_X = k, \mathbf{Z}' \setminus \{Z_X\}) \\ &\quad + \underbrace{\sum_{C \in \text{ch}(S) \cap \mathbf{N}''} w_{S,C} C(\mathbf{X}, Z_X = k, \mathbf{Z}' \setminus \{Z_X\})}_{=0} \end{aligned} \quad (18)$$

$$= D^{X,k}(\mathbf{X}^D) \sum_{C \in \text{ch}(S) \cap \mathbf{N}'} w_{S,C} C(\mathbf{X}^{R'}, Z_X = k, \mathbf{Z}' \setminus \{Z_X\}), \quad (19)$$

i.e. (16) holds for sum nodes. When \mathbf{N}_{k+1} is a product P, it must have a single child C in \mathbf{N}' , since P itself is contained in \mathbf{N}' . All other children are neither in \mathbf{N}' nor \mathbf{N}'' , due to decomposability. It is immediate, that when (16) holds for C, it also holds for P. Thus, (16) holds for all \mathbf{N}' , in particular for the root \mathbf{N}_K , and the proposition follows. \square