

# Sum-Product Network Decompileation

**Cory J. Butz**

*University of Regina, Regina, Canada*

BUTZ@CS.UREGINA.CA

**Jhonatan S. Oliveira**

*University of Regina, Regina, Canada*

OLIVEIRA@CS.UREGINA.CA

**Robert Peharz**

*Eindhoven University of Technology, Eindhoven, Netherlands*  
*University of Cambridge, Cambridge, United Kingdom*

RP587@CAM.AC.UK

## Abstract

There exists a dichotomy between classical probabilistic graphical models, such as *Bayesian networks* (BNs), and modern tractable models, such as *sum-product networks* (SPNs). The former generally have intractable inference, but provide a high level of interpretability, while the latter admit a wide range of tractable inference routines, but are typically harder to interpret. Due to this dichotomy, tools to convert between BNs and SPNs are desirable. While one direction – compiling BNs into SPNs – is well discussed in Darwiche’s seminal work on *arithmetic circuit* compilation, the converse direction – *decompiling* SPNs into BNs – has received surprisingly little attention. In this paper, we fill this gap by proposing SPN2BN, an algorithm that decompiles an SPN into a BN. SPN2BN has several salient features when compared to the only other two works decompiling SPNs. Most significantly, the BNs returned by SPN2BN are *minimal* independence-maps that are more parsimonious with respect to the introduction of *latent* variables. Secondly, the output BN produced by SPN2BN can be precisely characterized with respect to a compiled BN. More specifically, a certain set of directed edges will be added to the input BN, giving what we will call the *moral-closure*. Lastly, it is established that our compilation-decompilation process is idempotent. This has practical significance as it limits the size of the decompiled SPN.

**Keywords:** Probabilistic graphical models; Sum-product networks; Bayesian networks

## 1. Introduction

There exists a trade-off between classical probabilistic graphical models and recent tractable probabilistic models. Classical models, such as *Bayesian networks* (BNs) (Pearl, 1988), provide high-level interpretability as conditional independence assumptions are directly reflected in the underlying graphical structure. However, the downside is that performing exact inference in BNs is NP-hard. In contrast, modern tractable probabilistic models, such as *sum-product networks* (Darwiche, 2009; Poon and Domingos, 2011), allow a wide range of tractable inference, but are harder to interpret. In order to combine advantages of both BNs and SPNs – which are complementary regarding interpretability and inference efficiency – tools to convert back and forth between these types of models are vital.

The direction compiling BNs into SPNs is well understood due to Darwiche’s work on BN compilation into *arithmetic circuits* (ACs) (Darwiche, 2009).<sup>1</sup> Since inference in ACs and SPNs can be performed in linear time of the network size, compilation amounts to finding an inference machine

---

1. ACs and SPNs are equivalent models. *Deterministic* models are typically referred to as ACs, while *non-deterministic* models are called SPNs. See Section 2 for further details.

with minimal inference cost. ACs can also take advantage of context-specific-independence in the BN parameters to further reduce the size of the AC.

The converse direction of SPN decompilation into BNs has received limited attention. This lack of attention can be understood historically. Since an original purpose of ACs was to serve as efficient inference machine for a *known* BN, decompilation would seem like a mere academic exercise. The proposition of SPNs, however, introduced some practical changes to the AC model. First, unlike ACs, SPNs are typically learned directly from data, i.e., a reference BN is not available. Thus, providing a corresponding BN would greatly improve the interpretability of the learned SPN. Second, as already mentioned, SPNs are typically non-deterministic, which naturally introduces an interpretation of SPNs as hierarchical latent variable models (Peharz et al., 2017; Choi and Darwiche, 2017). A decompilation algorithm for SPNs should account for this fact, and generate BNs with a plausible set of latent variables. Thus, naively decompiling SPNs with a decompilation algorithm devised for ACs would yield densely connected and rather uninterpretable BNs.

In this paper, we fill this void by formalizing SPN decompilation. We propose SPN2BN, an algorithm that converts a trained SPN into a BN. Our algorithm arguably improves over the only two other approaches in the literature addressing the connection between BNs and SPNs (Zhao et al., 2015) and (Peharz et al., 2017). First, while (Peharz et al., 2017) produce a BN for a given SPN, these BNs are, in general, not minimal *independence-maps* (I-maps) (Pearl, 1988), i.e., they introduce needless dependence assumptions. Our algorithm SPN2BN, on the other hand, produces *minimal* I-maps. Second, (Zhao et al., 2015) is excessive with the number of introduced latent variables. In fact, both approaches interpret each single sum node in an SPN as a latent variable on its own. In this paper, we devise a more economical approach and identify groups of sum nodes to jointly represent one latent variable. This grouping is based on whether sum nodes are “on the same level of circuit hierarchy” and “responsible” for the same set of observable variables (these notions will be made formal in Section 3).

These design choices for SPN2BN improve over (Zhao et al., 2015) and (Peharz et al., 2017) both in terms of a reduced number of BN nodes (latent variables) and a reduced number of edges (minimal I-mapness). While this design leads to more succinct and perhaps more esthetic BNs, SPN2BN is also justified in a formal way. We show that SPN2BN can be seen as the inverse of the SPN compilation process proposed in (Darwiche, 2003). Consider an arbitrary BN  $\mathcal{B}$  that was compiled into an AC with *variable elimination* following a *reverse topological order* (Verto). Convert the AC into an SPN  $\mathcal{S}$  with an optional marginalization operation<sup>2</sup> (Darwiche, 2009) creating latent variables by rendering random variables unobserved. Then, decompiling SPN  $\mathcal{S}$  with SPN2BN yields a BN  $\mathcal{B}^c$  with a set of directed edges that are a superset of the original BN, which we call the *moral closure* of  $\mathcal{B}$ . The SPN2BN algorithm is consistent with respect to this compilation algorithm, in the sense that it always yields the moral closure  $\mathcal{B}^c$  of any given BN  $\mathcal{B}$ . Repeating the compilation-decompilation process with the morally closed BN is then idempotent, i.e. compilation and decompilation are consistent inverses of each other. Consistency with a compilation procedure is arguably a desirable property as it serves as a characterization of the decompilation method. In contrast, (Zhao et al., 2015) and (Peharz et al., 2017) are not consistent with any general-purpose compilation algorithm, and tend to excessively increase the number of variables and edges in the constructed BN. Lastly, even when the input SPN does not stem from an assumed compiler, e.g.,

---

2. SPNs are closed under marginalization, that is, any sub-marginal of any SPN can again be represented as an SPN.




Figure 1: Compilation of the BN in (i) using VErto (Darwiche, 2009), and marginalizing  $H_1$ ,  $H_2$ , and  $H_3$ , yields the SPN in (ii).

when it is learned from data, the VErto compilation assumption within SPN2BN helps us to interpret the result of decompilation.

For example, consider a prominent example of a BN in Figure 1 (i), commonly known as a hidden Markov model (HMM). In Figure 1 (ii), we see the result of VErto, followed by marginalization of  $H_1$ ,  $H_2$ , and  $H_3$  deeming these three variables latent. This SPN shall be converted back into a BN. In Figure 2 (i, ii) we see the BNs produced by (Zhao et al., 2015) and (Peharz et al., 2017), respectively. Both BNs introduce more variables than were present originally, and the introduced edges hardly reflect the succinct independence assumptions of the HMM. In Figure 2 (iii), the decompilation result by our SPN2BN algorithm is depicted. It can be seen that SPN2BN recovers the original HMM structure, where latent variables  $Z_1$ ,  $Z_2$ , and  $Z_3$  exactly correspond to the original latent variables  $H_1$ ,  $H_2$  and  $H_3$ , respectively. Evidently, no decompilation is able to recover the original labels for these variables, since reference to these has been explicitly removed by the previous (optional) marginalization operation. However, we see that SPN2BN successfully detects their signature in the compiled SPN, enabling it to recover an equivalent set of latent variables.

## 2. Sum-Product Networks

Here, we review BNs, ACs, and SPNs, as well as the compilation of BNs into SPNs.

We denote *random variables* (RVs) by uppercase letters, such as  $X$  and  $Y$ , possibly with subscripts, and their values by corresponding lowercase letters  $x$  and  $y$ . Sets of RVs are denoted by boldfaced uppercase letters and their combined values by corresponding boldfaced lowercase letters. The *children* of a variable  $V$  in a *directed acyclic graph* (DAG)  $\mathcal{B}$ , denoted  $Ch(V)$ , are the




Figure 2: Decompile of the SPN in Figure 1 (ii) by (Zhao et al., 2015) in (i), (Peharz et al., 2017) in (ii), and SPN2BN in (iii)

immediate descendants of  $V$  in  $\mathcal{B}$ . Similarly, the *parents*  $Pa(V)$  of a variable  $V$  are immediate ancestors of  $V$  in  $\mathcal{B}$ . The *descendants*  $De(V)$  are the variables  $V'$  with a directed path from  $V$  to  $V'$  in  $\mathcal{B}$ . The *ancestors*  $An(V)$  of  $V$  are similarly defined. A variable  $V_k$  is called a *v-structure* in a DAG  $\mathcal{B}$ , if directed edges  $(V_i, V_k)$  and  $(V_j, V_k)$  appear in  $\mathcal{B}$ , where  $V_i$  and  $V_j$  are non-adjacent variables in  $\mathcal{B}$ .

The independency information encoded in a DAG can be read graphically by the *d-separation* algorithm in linear time (Geiger et al., 1989).

**Definition 1** (Pearl, 1988) If  $\mathbf{X}$ ,  $\mathbf{Y}$ , and  $\mathbf{Z}$  are three disjoint subsets of nodes in a DAG  $\mathcal{B}$ , then  $\mathbf{Z}$  is said to *d-separate*  $\mathbf{X}$  from  $\mathbf{Y}$ , denoted  $I(\mathbf{X}, \mathbf{Z}, \mathbf{Y})$ , if along every path between a node in  $\mathbf{X}$  and a node in  $\mathbf{Y}$  there is a node  $W$  satisfying one of the following two conditions: (i)  $W$  is a v-structure and neither  $W$  nor any of its descendants are in  $\mathbf{Z}$ , or (ii)  $W$  is not a v-structure and  $W$  is in  $\mathbf{Z}$ .

The next definition formalizes when a DAG is an I-map of a joint probability distribution (JPD).

**Definition 2** (Darwiche, 2009) Let  $\mathcal{B}$  be a DAG and  $P$  be a JPD over the same set of variables.  $\mathcal{B}$  is an I-map of  $P$  if every conditional independence read by d-separation on  $\mathcal{B}$  holds in the distribution  $P$ . An I-map  $\mathcal{B}$  is *minimal*, if  $\mathcal{B}$  ceases to be an I-map when we delete any edge from  $\mathcal{B}$ .

BNs are DAGs with nodes representing variables and edges representing variable dependencies, in which the strength of these relationships are quantified by conditional probability tables (CPTs). More formally, a BN over variables  $\mathbf{X}$  has its CPTs defined over each variable given its parents, that is,  $P(V|Pa(V))$ , for every  $V \in \mathbf{X}$ . One salient feature is that the product of the BN CPTs yields a JPD  $P$  over  $\mathbf{X}$ .

**Definition 3** (Pearl, 1988) Given a JPD  $P$  on a set of variables  $\mathbf{X}$ , a DAG  $\mathcal{B}$  is called a *Bayesian network (BN)* of  $P$  if  $\mathcal{B}$  is a minimal I-map of  $P$ .

In a BN, the independencies read by d-separation in the DAG  $\mathcal{B}$  are guaranteed to hold in the JPD  $P$ . Unfortunately, while BNs have clear interpretability, exact inference in BNs is NP-hard.

BNs can be compiled into *Arithmetic Circuits* (ACs) (Darwiche, 2003) by graphically mapping the operations performed when marginalizing all variables from the BN.

**Definition 4** (Darwiche, 2003) *An arithmetic circuit (AC) over variables  $\mathbf{U}$  is a rooted DAG whose leaf nodes are labeled with numeric constants, called parameters, or  $\lambda$  variables, called indicators, and whose other nodes are labeled with multiplication and addition operations.*

Notice that parameter variables are set according to the BN CPTs, while indicator variables are set according to any observed evidence.

SPNs are a probabilistic graphical model that can be learned from data using, for instance, the LearnSPN algorithm (Gens and Domingos, 2013).

**Definition 5** (Poon and Domingos, 2011) *A sum-product network (SPN) is a DAG containing three types of nodes: leaf distributions, sums, and products. Leaves are tractable distribution functions over  $\mathbf{Y} \subseteq \mathbf{X}$ . Sum nodes  $S$  compute weighted sums  $S = \sum_{N \in Ch(S)} w_{S,N} N$ , where  $Ch(S)$  are the children of  $S$  and  $w_{S,N}$  are weights that are assumed to be non-negative and normalized (Peharz et al., 2015). Product nodes  $P$  compute  $P = \prod_{N \in Ch(P)} N$ . The value of an SPN, denoted  $\mathcal{S}(\mathbf{x})$ , is the value of its root.*

The *scope* of a sum or product node  $N$  is recursively defined as  $sc(N) = \bigcup_{C \in Ch(N)} sc(C)$ , while the scope of a leaf distribution is the set of variables over which the distribution is defined. A *valid* SPN defines a JPD and allows for efficient inference (Poon and Domingos, 2011). The following two structural constraints on the DAG guarantee validity. An SPN is *complete* if, for every sum node, its children have the same scope. An SPN is *decomposable* if, for every product node, the scopes of its children are pairwise disjoint. Valid SPNs are of particular interest because they represent a JPD over the variables in the problem domain. In addition, like ACs, exact inference is linear in the size of the DAG. Unlike ACs, however, SPNs allow for a latent variable (LV) interpretation (Peharz et al., 2017).

In (Poon and Domingos, 2011), it was suggested that SPNs can be interpreted as hierarchical latent variable model, where each sum node corresponds to a latent, marginalized random variable. This interpretation can be made explicit, by incorporating the latent variables *explicitly* in the SPN, yielding the so-called augmented SPN (Peharz et al., 2017). We briefly review the construction of the augmented SPN here: first, for each sum node  $S$  we postulate a random variable  $Z_S$  with  $|Ch(S)|$  states, i.e. each state of  $Z_S$  corresponds to one of  $S$ 's children. The states of  $Z_S$  are represented via indicators  $\lambda_{Z_S=k}$ , which are explicitly introduced in the SPN. Furthermore, the sum node  $S = \sum_k w_{S,k} C_k$  is replaced with  $S = \sum_k w_{S,k} C_k * \lambda_{Z_S=k}$ . This construction, originally proposed by (Poon and Domingos, 2011), allows us to switch  $S$ 's children “off and on” by setting the indicators of  $Z_S$ , and therefore interpret the children as distributions, conditioned on  $Z_S$ .

However, as observed in (Peharz et al., 2017), this construction is in conflict with the completeness requirement of SPNs. Consider a sum node  $S^c \in An(S)$  with the property that it has a child  $C \in Ch(S^c)$  which does not reach  $S$ , i.e.,  $S \notin De(C)$ . (Peharz et al., 2017) call such  $S^c$  a *conditioning sum* of  $S$ . Including the indicators as above renders  $S^c$  incomplete, since some but not all children of  $S^c$  reach  $Z_S$ . This has the severe consequence that the tractable inference mechanism for SPNs is now invalid. (Peharz et al., 2017) propose to fix this problem by introducing a new (dummy) sum node  $\bar{S}$ , which has only the indicators  $\lambda_{Z_S=k}$ ,  $k = 1 \dots, |Ch(S)|$  as children, the so-called *twin sum* of  $S$ . Each child  $C \in Ch(S^c)$  which does not reach  $S$  is now replaced with  $C * \bar{S}$ .

---

**Algorithm 1** SPN Compilation Assumption

---

```
1: BN2SPN( $\mathcal{B}$ )
2: Let  $\sigma$  be a reverse topological ordering of  $\mathcal{B}$       ▷ VE with reverse topological order (VERto)
3:  $\mathcal{C}$  = compile-to-AC-with-VE( $\mathcal{B}, \sigma$ )
4:  $\mathcal{S}$  = redistribute-parameters( $\mathcal{C}$ )                      ▷ Convert AC to SPN
5:  $\mathcal{S}$  = compile-marginalized-spn( $\mathcal{S}$ )
6: changed = true
7: while changed do
8:   Let  $\mathcal{S}'$  be a copy of  $\mathcal{S}$ 
9:    $\mathcal{S}$  = add-terminal-nodes( $\mathcal{S}$ )
10:   $\mathcal{S}$  = remove-products-of-products( $\mathcal{S}$ )
11:   $\mathcal{S}$  = lump-products( $\mathcal{S}$ )                               ▷ Lump products over the same children
12:  if  $\mathcal{S} == \mathcal{S}'$  then
13:    changed = false
14: Return  $\mathcal{S}$ 
```

---

As shown in (Peharz et al., 2017), this process leads to a consistent augmentation of the SPN in that sense that it explicitly manifests a random variable for  $S$ , while i) maintaining completeness and decomposability, and ii) leaving the marginal distribution over observed variables  $X$  unchanged. For further details, see (Peharz et al., 2017), in particular Algorithm AugmentSPN.

### 3. SPN Decompilation

In this section, we formalize SPN decompilation into a BN.

The interpretation of SPNs as latent variable models is not unique. For instance, every SPN sum node can be viewed itself as a LV, as done in (Zhao et al., 2015; Peharz et al., 2015). In stark contrast, all SPN sum nodes can be interpreted as one single LV (Peharz, 2015). This provides a wide spectrum of interpretations based only on those sum nodes appearing in an SPN. In addition, external LVs can be introduced to an SPN, such as the *switching parents* in (Peharz et al., 2015). Thus, for a given SPN learned from data, there are seemingly countless possible interpretations of its latent space.

The approach taken in this paper is to make a *compilation assumption*, i.e., to assume that there exist some underlying BN which was compiled into the SPN at hand. While there are many possible ways to compile a BN into an SPN, we use arguably the most prominent compilation method, *Variable Elimination* (VE) (Zhang and Poole, 1994) following any *reverse topological order* (VERto). In particular, the recursive marginalization of variables during VE generate hierarchical layers of sum nodes, which typically appear in SPNs. Besides assuming that the underlying SPN was generated from a BN, we further assume that some of the BN’s variables have been removed, that is, marginalized from the model. The task of decompilation can then be formulated to recover the original BN structure as far as possible. The main contributions in this paper are to i) provide such an algorithm SPN2BN, and ii) show that it indeed recovers the *morally closed* version of the original BN (see Definition 11 for moral closure).

Algorithm 1 describes our compilation assumption, VERto followed by optionally marginalizing some variables. In line 3, a given BN  $\mathcal{B}$  is converted into an AC  $\mathcal{C}$  using VERto (Darwiche, 2003). In

line 4, the leaf parameters are redistributed as sum-weights (Rooshenas and Lowd, 2014), yielding an SPN  $\mathcal{S}$ . In line 5, we assume that all internal latent variables in  $\mathcal{S}$  are marginalized and, thus, all of their indicator variables are set to 1. Here, any arbitrary subset of the internal latent variables can be considered. Next, we recursively simplify  $\mathcal{S}$  by applying three operations until no further change can be made. In line 9, a sum node with only indicator nodes as children is converted into a *terminal node* (Zhao et al., 2015), which is a univariate distribution over the indicator variable. Product nodes whose children are exclusively products, i.e., chains of products are simplified into a single product node in line 10. Finally, in line 12, if two or more product nodes have the same set of children, then they are lumped into a single product node.

On the other hand, by *decompilation*, we mean the procedure of converting an SPN into a BN. This process involves determining the RVs and DAG for the BN. We can suggest RVs for the BN by analyzing the compilation assumption. Similarly, an I-map can be obtained as a DAG using the SPN DAG. We now formalize these ideas.

**Definition 6** *Given an SPN over RVs  $\mathbf{X}$  and a compilation assumption, SPN decompilation is an algorithm that both: (i) suggests a set of LVs  $\mathbf{Z}$ , and; (ii) produces an I-map over  $\mathbf{X}$  and  $\mathbf{Z}$ .*

Task (i) of SPN decompilation is more involved than expected. A naive approach is to disregard the compilation assumption and treat each sum node as one LV. Negative consequences of this approach will be discussed in the next section. We suggest a more elegant approach by interpreting the effect of the compilation assumption on the graphical characteristics of the SPN.

Recall that we assume the SPN was compiled using Verto. During compilation, marginalizing variables creates groups of sum nodes in the same layer (the distance of the longest path from the root). Hence, identifying these groups is a way of suggesting RVs for the decompiled BN.

More formally, given a sum node  $S$ , the *sum-depth* of  $S$  is the number of sum nodes in the longest directed path from the root to  $S$ .

**Example 1** *The sum-depth of sum node  $S_3$  in the SPN of Figure 1 (ii) is 2, since there are 2 sum nodes on the longest path from the root to  $S_3$ . Similarly, the sum-depth of  $S_2$  is 1 and of  $S_1$  is 0.*

A *sum-layer* is the set of all sum nodes having the same sum-depth.

**Example 2** *One sum-layer in the SPN of Figure 1 (ii) consists of  $S_3$  and  $S_5$ , since both  $S_3$  and  $S_5$  have a sum-depth of 2. Furthermore,  $S_2$  and  $S_4$  form another sum-layer, as does  $S_1$  by itself.*

A *sum-region* is the set of all sum-nodes within the same sum-layer and having the same scope.

**Example 3** *Sum-layer  $S_3$  and  $S_5$  in the SPN of Figure 1 (ii) has only one sum-region, since  $S_3$  and  $S_5$  have the same scope. For the same reason, sum-layer  $S_2$  and  $S_4$  also has only one sum-region.*

A sum-region is created by marginalizing variables during our compilation assumption. Thus, to answer task (i) of SPN decompilation, we suggest that  $\mathbf{Z}$  consists of one LV per sum-region.

**Example 4** *In the SPN of Figure 1 (ii), we suggest three LVs in  $\mathbf{Z} = \{Z_{S_1}, Z_{S_2}, Z_{S_3}\}$ , namely, one per sum-region.*

---

**Algorithm 2** SPN Decompileation

---

```
1: SPN2BN( $\mathcal{S}$ )
2: Let  $L$  be the list of sum-layers in  $\mathcal{S}$ 
3: Let  $\mathbf{Z}$  denote a mapping from nodes to LVs
4: for each node  $N$  in  $\mathcal{S}$  do
5:    $\mathbf{Z}[N] = \emptyset$ 
6:                                      $\triangleright$  Phase (i) suggests a set  $\mathbf{Z}$  of LVs
7:   for each  $l$  in  $L$  do                                      $\triangleright$  for each layer
8:      $\mathbf{S} = \emptyset$                                           $\triangleright$  Initialization of Scopes
9:     for each node  $N$  in  $l$  do                                $\triangleright$  for each node
10:      Let  $X$  denote  $\text{scope}(N)$ 
11:      if  $X \in \mathbf{S}$  then
12:        Let  $Z$  be the existing LV for scope  $X$ 
13:         $\mathbf{Z}[N] = Z$ 
14:      else
15:        Let  $Z'$  be a new LV
16:         $\mathbf{Z}[N] = Z'$                                         $\triangleright$  Update scopes
17:         $\mathbf{S} = \mathbf{S} \cup \{X\}$ 
18:                                      $\triangleright$  Phase (ii) produces an I-map over  $\mathbf{X}$  and  $\mathbf{Z}$ 
19:   for each node  $N$  in  $\mathcal{S}$  do
20:     if  $N$  is a sum or leaf node then
21:       for each sum node  $S'$  in  $An(N)$  do
22:         if  $S'$  is conditioning w.r.t.  $N$  then
23:           Add edge  $(\mathbf{Z}[S'], N)$  to  $\mathcal{B}$ 
24:   Return  $\mathcal{B}$ 
```

---

We now turn our attention to task (ii) of SPN decompileation, that is, constructing an I-map over  $\mathbf{X}$  and  $\mathbf{Z}$ . Augment the SPN as done in (Peharz et al., 2017). However, before continuing, we need to correct the notion of a conditioning sum node for the following reason. Consider sum node  $S_3$  in the SPN of Figure 1 (ii). (Peharz et al., 2017) would not define sum node  $S_1$  as a conditioning sum node for  $S_3$ , even though  $Z_{S_1}$  would appear as a conditioning variable for  $Z_{S_3}$  in the CPT  $P(Z_{S_3}|Z_{S_1}, Z_{S_2}, Z_{S_4})$ , as depicted in the constructed I-map in Figure 2 (ii).

**Definition 7** *An ancestor sum node  $S$  of a node  $N$  in an augmented SPN is called conditioning, if it is not true that all children of  $S$  reach exactly the same subset of  $S$  and  $\bar{S}$ .*

**Example 5** *Consider sum node  $S_3$  in the SPN of Figure 1 (ii). Ancestor sum node  $S_2$  is conditioning, since the left-most child of  $S_2$  reaches  $S_3$ , but the right-most child does not. Node  $S_1$  is not conditioning for  $S_3$ , since all children of  $S_1$  reach the same subset  $S_3$  and  $\bar{S}_3$  in the augmented SPN.*

In Example 5, observe that  $S_1$  is not a conditioning sum node for  $S_3$  and hence  $Z_{S_1}$  does not appear as parent of  $Z_{S_3}$  in our constructed I-map in Figure 2 (iii).

The SPN decompileation techniques described thus far are formalized as Algorithm 2.



**Example 6** Given the SPN in Figure 1 (ii) as input to  $\text{SPN2BN}(S)$ , the decompiled BN is given in Figure 2 (iii).

## 4. Theoretical Foundation

In this section, we first establish important properties of SPN decompilation. Later, we show a favorable characteristic of our compilation assumption and Algorithm 2.

### 4.1 On SPN Decompilation

Our decompilation algorithm is parsimonious with the introduction of LVs. One LV is assigned per sum-region rather than one per sum node.

Regarding I-map construction, we first show the correctness of the I-map, and then establish that the constructed I-map is minimal.

The I-map correctness follows from the CPT construction suggested in (Peharz et al., 2017). Theorem 1 in (Peharz et al., 2017) shows that certain independencies necessarily hold in an SPN, namely, each LV  $Z_S$  for a sum node  $S$  is conditionally independent of all non-descendant sum nodes given all ancestor sum nodes of  $S$ , denoted  $\mathbf{Z}_P$ . A CPT  $P(Z_S|\mathbf{Z}_P)$  is constructed for their I-map. The I-map built in Algorithm 2 uses the same CPT probability values, except building the CPT  $P(Z_S|\mathbf{Z}_C)$ , where  $\mathbf{Z}_C$  are those conditioning nodes defined in Definition 7. Since  $\mathbf{Z}_C \subseteq \mathbf{Z}_P$ , the independencies encoded in the I-map of (Peharz et al., 2017) are a subset of those encoded by the I-map built by Algorithm 2. Thus, the I-maps proposed in (Peharz et al., 2017) are, in general, not minimal.

**Example 7** Consider sum node  $S_3$  in the SPN of Figure 1 (ii). For (Peharz et al., 2017), the CPT for  $Z_{S_3}$  is  $P(Z_{S_3}|Z_{S_1}, Z_{S_2}, Z_{S_4})$ . Here,  $Z_{S_3}$  is independent of  $Z_{S_1}$ , given  $Z_{S_2}$  and  $Z_{S_4}$ . Thus, Algorithm 2 builds the smaller CPT  $P(Z_{S_3}|Z_{S_2}, Z_{S_4})$  for  $Z_{S_3}$ .

The proof for these new conditional independencies and the correctness of our I-map is formalized in Lemma 8.

**Lemma 8** Consider an augmented SPN  $S$  with a sum node  $S$ . Let  $\mathbf{A}$  be all of  $S$ 's ancestors and  $\mathbf{C}$  all of  $S$ 's conditioning sum nodes. Then,  $P(Z_S|\mathbf{Z}_A) = P(Z_S|\mathbf{Z}_C)$ .

**Proof** Let  $\mathbf{N} = \mathbf{A} - \mathbf{C}$  be the non-conditioning ancestor sum nodes of  $S$ . While conditioning on  $\mathbf{Z}_A$  selects a single path in the augmented SPN from the root to  $S$ , multiple paths may exist from the root to  $S$  when conditioning only on  $\mathbf{Z}_C$ . By Definition 7, however, all children of a non-conditioning node reach the same subset of  $S$  and  $\bar{S}$ . Therefore, the weight  $w$  corresponding to the instantiation of  $Z_S$  necessarily appears in every term of products in the summation. By the distributive law,  $w$  can be pulled out of this summation of products. The resulting summation of products is precisely the SPN computation for the probability of the conditioning event. Hence, these two summation of products cancel each other out leaving the conditional probability of  $P(Z_S|\mathbf{Z}_C)$  to be  $w$ . Thus,  $P(Z_S|\mathbf{Z}_A) = P(Z_S|\mathbf{Z}_C)$ . ■

We next show that our constructed I-maps are minimal.

**Theorem 9** *Let  $S$  be an augmented SPN over RVs  $\mathbf{X}$  and LVs  $\mathbf{Z}$ . Given  $S$  as input, the algorithm SPN2BN builds a minimal I-map.*

**Proof** First, notice that, by Lemma 8, the output  $\mathcal{B}$  of SPN2BN is an I-map. By contradiction, suppose that the constructed I-map  $\mathcal{B}$  is not minimal. Then there exists a directed edge  $(S_C, S_N)$  from a conditioning node  $S_C$  for a sum node  $S_N$  that can be deleted without destroying I-mapness. In particular, this means that

$$P(S_N | Pa(S_N)) = P(S_N | Pa(S_N) - S_C). \quad (1)$$

By Definition 7, as  $S_C$  is a conditioning sum for  $S_N$ , there exist at least two children  $S_i$  and  $S_j$  of  $S_N$  that select different paths to  $S_N$  and its twin  $\bar{S}_N$ . Since the respective weights  $w_i$  and  $w_j$  of  $S_i$  and  $S_j$  can be different, it immediately follows that Equation (1) is not satisfied by the joint probability distribution  $P(\mathbf{XZ})$  defined by  $S$ . Thus, the I-mapness is violated, if  $(S_C, S_N)$  is removed from  $\mathcal{B}$ . Therefore, by contradiction,  $\mathcal{B}$  is a minimal I-map. ■

One seeks minimal I-maps as non-minimal I-maps are not necessarily useful in practice (Darwiche, 2009; Pearl, 1988; Koller and Friedman, 2009).

## 4.2 Compilation and Decompilation

In this section, we first show that BN2SPN2BN, our compilation-decompilation algorithm, constructs a unique BN for a given set of original BNs. A consequence of this is that BN2SPN2BN is idempotent.

We next show that the BN output by BN2SPN2BN can be different than the original BN. In the remainder of this section, we assume  $\prec$  is a fixed topological ordering of a given BN  $\mathcal{B}$ .

**Example 8** *Consider the call  $\text{BN2SPN2BN}(\mathcal{B})$ , where BN  $\mathcal{B}$  has directed edges  $\{(A, B), (B, E), (C, D), (D, E)\}$ . Then, the output BN has directed edges  $\{(A, B), (B, E), (C, D), (D, E), (B, D), (B, C)\}$ .*

Notice that the directed edges of the original BN are a subset of those in the output BN.

**Definition 10** *A directed moralization edge is a directed edge  $(V_i, V_j)$  added between two non-adjacent vertices  $V_i$  and  $V_j$  in a given BN  $\mathcal{B}$  whenever there exists a variable  $V \in \mathcal{B}$  such that  $V \in \text{Ch}(V_i)$  and  $V \in \text{Ch}(V_j)$ , where  $V_i \prec V_j$ .*

We now introduce the key notion of moral closure.

**Definition 11** *Given a BN  $\mathcal{B}$  and a fixed topological order  $\prec$  of  $\mathcal{B}$ , the moral closure of  $\mathcal{B}$ , denoted  $\mathcal{B}^c$ , is the unique BN formed by iteratively augmenting  $\mathcal{B}$  with all directed moralization edges.*

We are now ready to present the first main result of our compilation-decompilation process.

**Theorem 12** *Given a BN  $\mathcal{B}$  and a fixed topological order  $\prec$  of  $\mathcal{B}$ , the output of the compilation-decompilation algorithm BN2SPN2BN is the moral closure  $\mathcal{B}^c$  of  $\mathcal{B}$ .*

**Proof** Applying BN2SPN2BN on  $\mathcal{B}$  involves running BN2SPN followed by SPN2BN. Consider running BN2SPN on  $\mathcal{B}$ . This involves eliminating all variables from  $\mathcal{B}$  following an elimination ordering  $\sigma$ . It is well-known that this process builds a triangulated, undirected graph (Pearl, 1988). Triangulated, undirected graphs admit a perfect numbering. This means that no fill-in edges need to be added to the triangulated, undirected graph when eliminating variables following the perfect numbering.

Now  $\sigma$  itself is a perfect numbering for the triangulated graph built by eliminating variable following  $\sigma$ . Let us focus on the undirected edges that were added to  $\mathcal{B}$ . There are two cases to consider, i.e., moralization edges and triangulation edges. A moralization edge corresponds to a directed moralization edge in our case.

A triangulation edge  $(V_i, V_j)$  is added if and only if  $V_i$  and  $V_j$  are non-adjacent parents of a common child  $V_k$  in the directed case. However, these edges are precisely the directed moralization edges that are recursively added. Therefore, BN2SPN builds an SPN following the hierarchy in the moral closure of  $\mathcal{B}$  and subsequently SPN2BN unwinds the SPN following this same hierarchy. Thus, given  $\mathcal{B}$ , BN2SPN2BN yields the moral closure  $\mathcal{B}^c$  of  $\mathcal{B}$ . ■

Theorem 12 has a couple of important consequences. As Theorem 12 establishes that the output of BN2SPN2BN is the moral closure  $\mathcal{B}^c$  of the input BN  $\mathcal{B}$ , it immediately follows that the output BN is exactly the input BN whenever no directed moralization edges are added to  $\mathcal{B}$ . One situation where this occurs is when  $\mathcal{B}$  does not have any v-structures, such as in the case of HMMs. Here,  $\mathcal{B}^c = \mathcal{B}$ , so the output BN of BN2SPN2BN is the same as the input (up to a relabelling of variables). For example, recall the HMM in Figure 1 (i). BN2SPN2BN yielded back the same BN as illustrated in Figure 2 (iii). A second important case is when the input is  $\mathcal{B}^c$  itself. This leads to our next result showing that our compilation-decompilation process is idempotent.

**Theorem 13** *BN2SPN2BN is idempotent.*

**Proof** Let  $\mathcal{B}$  be a BN and  $\prec$  a fixed topological ordering of  $\mathcal{B}$ . Then  $\mathcal{B}^c = \text{BN2SPN2BN}(\mathcal{B})$ , by Theorem 12. By definition, the moral closure of  $\mathcal{B}^c$  is  $\mathcal{B}^c$  itself. Thus,  $\mathcal{B}^c = \text{BN2SPN2BN}(\mathcal{B}^c)$ . Therefore, BN2SPN2BN is idempotent. ■

Theorem 13 has practical significance because it limits the maximum size of the decompiled BN to be the size of the moral closure of the input BN. In contrast, if we change the decompilation method to (Zhao et al., 2015) or (Peharz et al., 2017), then applying BN2SPN2BN repeatedly will continually yield a larger BN.

## 5. Conclusion

In this paper, we formalize SPN decompilation by suggesting SPN2BN, an algorithm that converts an SPN into a BN. SPN2BN is an improvement over (Zhao et al., 2015) and (Peharz et al., 2017), which are excessive with the number of introduced latent variables. One key result of our SPN decompilation is that it constructs the moral closure  $\mathcal{B}^c$  of the original BN  $\mathcal{B}$ . This means that in certain cases like for HMMs, where the moral closure of a BN  $\mathcal{B}$  is  $\mathcal{B}$  itself, our SPN decompilation will return the original BN. Moreover, our compilation-decompilation process is idempotent. This has practical significance as it limits the maximum size of the decompiled BN to be the size of  $\mathcal{B}^c$ .

## Acknowledgement

This research is partially supported by NSERC Discovery Grant 238880 and has also received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 797223 — HYBSPN.

## References

- A. Choi and A. Darwiche. On relaxing determinism in arithmetic circuits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 825–833, 2017.
- A. Darwiche. A differential approach to inference in Bayesian networks. *Journal of the ACM*, 50(3):280–305, 2003.
- A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, Los Angeles, CA, 2009.
- D. Geiger, T. S. Verma, and J. Pearl. d-separation: From theorems to algorithms. In *Proceedings of the Fifth Conference on Uncertainty in Artificial Intelligence*, pages 139–148, 1989.
- R. Gens and P. Domingos. Learning the structure of sum-product networks. In *Proceedings of the Thirtieth International Conference on Machine Learning*, pages 873–880, 2013.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, 2009.
- J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco, CA, 1988.
- R. Peharz. *Foundations of sum-product networks for probabilistic modeling*. PhD thesis, 2015.
- R. Peharz, S. Tschiatschek, F. Pernkopf, and P. Domingos. On theoretical properties of sum-product networks. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 744–752, 2015.
- R. Peharz, R. Gens, F. Pernkopf, and P. Domingos. On the latent variable interpretation in sum-product networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(10):2030–2044, 2017.
- H. Poon and P. Domingos. Sum-product networks: A new deep architecture. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 337–346, 2011.
- A. Rooshenas and D. Lowd. Learning sum-product networks with direct and indirect variable interactions. In *Proceedings of the Thirty-First International Conference on Machine Learning*, pages 710–718, 2014.
- N. L. Zhang and D. Poole. A simple approach to Bayesian network computations. In *Proceedings of the Tenth Canadian Artificial Intelligence Conference*, pages 171–178, 1994.
- H. Zhao, M. Melibari, and P. Poupart. On the relationship between sum-product networks and Bayesian networks. In *Proceedings of Thirty-Second International Conference on Machine Learning*, pages 116–124, 2015.