
Structural Learning of Probabilistic Sentential Decision Diagrams under Partial Closed-World Assumption

Alessandro Antonucci¹

Alessandro Facchini*¹

Lilith Mattei¹

¹Istituto Dalle Molle di Studi per l'Intelligenza Artificiale (IDSIA), Lugano, Switzerland

Abstract

Probabilistic sentential decision diagrams are a class of structured-decomposable probabilistic circuits especially designed to embed logical constraints. To adapt the classical LEARNSPN scheme to learn the structure of these models, we propose a new scheme based on a partial closed-world assumption: data implicitly provide the logical base of the circuit. Sum nodes are thus learned by recursively clustering batches in the initial data base, while the partitioning of the variables obeys a given input vtree. Preliminary experiments show that the proposed approach might properly fit training data, and generalize well to test data, provided that these remain consistent with the underlying logical base, that is a relaxation of the training data base.

1 INTRODUCTION

Probabilistic Circuits (PCs) have been recently introduced as a general computational framework unifying the different existing formalisms for tractable probabilistic modeling [Choi et al., 2020].

Learning the structure of a PC from data is emerging as a crucial challenge for a widespread application of PCs in the area of machine learning. Different algorithms have been (e.g., Lowd and Domingos [2012]) and still are (e.g., Peharz et al. [2020]) designed to improve the state of the art in this field. Most of these algorithms are somehow inspired by the LEARNSPN scheme of Gens and Domingos [2013], the first algorithm for the structural learning of general PCs, originally designed for sum-product networks [Poon and Domingos, 2011], the most popular class of PCs.

In this position paper, we focus on the problem of learning from data, without other prior domain knowledge, the

structure of a particular class of PCs, called *Probabilistic Sentential Decision Diagrams* (PSDDs, Kisa et al. [2014]). Roughly speaking, a PSDD is a weighted logical circuit inducing a probability distribution assigning non-zero probability only to states consistent with the Boolean formula encoded by the circuit.

As a matter of fact there are only two main algorithms to learn the structure of a PSDD: LEARNPSDD [Liang et al., 2017] and STRUDEL [Dang et al., 2020]. Despite their differences, both these algorithms start from a circuit subject to local transformations. These local actions are designed to preserve the underlying logical base of the circuit, which should therefore be specified when the initial circuit is given, meaning that both algorithms may assume some prior domain knowledge modelled by the input circuit.

The LEARNSPN architecture is different: the algorithm learns a circuit from an input database by growing it from the root, rather than performing local operations on an initial circuit. LEARNSPN takes as input a database and recursively divides it into sub-databases by finding almost independent subsets of variables and clustering similar instances.

Two structural properties of PSDDs prevent a straightforward application of LEARNSPN. First, PSDDs are structured-decomposable, this roughly meaning that the splitting of the variables into independent sets must agree to some constraint. Second, a PSDD encodes a Boolean formula representing a knowledge base (KB); thus it is desirable to steer the operations of LEARNSPN so that the obtained KB represents the input database. Here we address such situation by making LEARNSPN able to learn a PSDD encoding a KB describing the data.

Of course there is no unique way to agree on what “describing the data” means. In database theory, a traditional assumption, called *closed-world assumption* (CWA), states that a database reflects with fidelity the concerned state of affairs, that it is complete [Reiter, 1981]. Stated otherwise, a fact that is true in the world is stored in the database, whereas a fact that is missing in the database just does not hold in the

*Supported by the Hasler foundation grant n. 20061.

world. This assumption is usually followed in ILP (Inductive Logic Programming), a branch of Machine Learning whose objective is to induce a set hypothesis represented as logic programs (rules). The opposite of the closed-world assumption is the *open-world assumption* [Reiter, 1981, Imieliński and Lipski Jr, 1989]. Commonly employed in the area of description logic [Baader et al., 2003], this simply states that some facts holding in the real world may be missing, and thus does not presume that the knowledge of a domain is complete.

However, in many cases, it may be desirable to have an intermediate perspective: some parts of the observed data may be complete, and some others may be incomplete. First treated by Motro [1989], this view has later been instantiated in various forms under various names (local closed-world assumption, partial closed-world assumption, partial completeness, and others; see, e.g., Galárraga et al. [2013], Darari et al. [2013], Dong et al. [2014])

In this work we follow a similar intermediate perspective. More precisely, we extract a KB from a database \mathcal{D} by recursively applying an instance of a *partial closed-world assumption* (PCWA). That is, each step of the procedure extracts a sub-database \mathcal{D}' and a subset \mathbf{X}' of variables for which \mathcal{D}' is assumed to reflect completely the concerned possible states of affairs (given the constraints provided by the previous step). This idea is made precise in Section 3.

The paper is organised as follows. In Section 2 we provide the necessary background about PSDDs. Our contribution is in Section 3, where we present a PCWA based structural learning algorithm for PSDDs inspired by the LEARNSPN schema. In Section 4 we run preliminary experiments. Conclusions and outlooks are in Section 5.

2 BACKGROUND

2.1 BASICS

We focus on a set of Boolean variables $\mathbf{X} := (X_1, \dots, X_n)$. Let \mathcal{D} denote a database of joint states of \mathbf{X} , to be also called *records* and denoted as $\mathbf{x} = (x_1, \dots, x_n)$. The database obtained by taking only the columns of \mathcal{D} relative to variables in $\mathbf{X}' \subseteq \mathbf{X}$ is denoted as $\mathcal{D}^{\mathbf{X}'}$. When coping with a single variable $X \in \mathbf{X}$, with a small abuse of notation, we intend as $\mathcal{D}^{X=1}$ the (sub)set of true instances of \mathcal{D}^X , and we analogously intend $\mathcal{D}^{X=0}$. Notation $|\mathcal{D}|$ is used instead to denote the number of records in \mathcal{D} . If the records are indexed by the integers $\{1, \dots, m\}$, with $m := |\mathcal{D}|$, we denote as \mathcal{D}_J the database obtained by taking only the records indexed by the elements of J , with $J \subseteq \{1, \dots, m\}$. Let $\mathbf{x}^j = x_1^j, \dots, x_n^j$ denote the j -th record of \mathcal{D} . According to CWA, we assume the database $\mathcal{D} = \{\mathbf{x}^1, \dots, \mathbf{x}^m\}$ complete, this meaning that it completely describes the possible state of (the concerned aspects of) the world. Hence, \mathcal{D}

induces a Disjunctive Normal Form (DNF)

$$\phi(\mathcal{D}) := \bigvee_{j=1}^m \bigwedge_{i=1}^n L_i^j, \quad (1)$$

whose conjunctive clauses are the ones describing the records belonging to \mathcal{D} . That is, for each record $\mathbf{x}^j := (x_1^j, \dots, x_n^j) \in \mathcal{D}$, the corresponding conjunctive clause under the CWA coincides with $L_1^j \wedge \dots \wedge L_n^j$ where, for $i \leq n$

$$L_i^j = \begin{cases} X_i & \text{when } x_i^j = 1, \\ \neg X_i & \text{otherwise.} \end{cases} \quad (2)$$

2.2 PSDDS

Inspired by Choi et al. [2020], in this section we introduce PSDDs as a class of PCs that, in addition to specific structural properties (e.g., determinism and structured-decomposability), encode Boolean formulae.

2.2.1 Probabilistic Circuits

A PC \mathcal{C} over \mathbf{X} is a rooted DAG \mathcal{G} annotated by parameters θ in order to represent a joint probability mass function $\mathbb{P}(\mathbf{X})$. Formally, each leaf (terminal node) g of \mathcal{G} is called an *input unit* and it is associated with a variable $X_g \in \mathbf{X}$ and a (possibly degenerate) probability mass function $d_g(X_g)$. Non-leaves nodes can be *sum units* and *product units*. Parameters are associated with both input units and sum units. In case of a sum unit, these can be seen as the probabilities of a mass function weighting its outgoing arcs, while for input units a mass function π is directly provided. In practice, if $\text{ch}(g)$ are the children of a node g (also called its *inputs*) and \mathbf{X}_g the set of variables associated to the input units reachable from g , and called its *scope*, the corresponding probability mass function $\mathbb{P}_g(\mathbf{X}_g)$ is defined as follows:

$$\mathbb{P}_g(\mathbf{x}_g) := \begin{cases} \pi_g(\mathbf{x}_g) & \text{if } g \text{ is an input unit,} \\ \prod_{\ell \in \text{ch}(g)} \mathbb{P}_\ell(\mathbf{x}_\ell) & \text{if } g \text{ is a product unit,} \\ \sum_{\ell \in \text{ch}(g)} \theta_{\ell,g} \mathbb{P}_\ell(\mathbf{x}_\ell) & \text{if } g \text{ is a sum unit.} \end{cases}$$

As we focus here on Boolean variables, we can identify input units with literals or logical constants, sum units with disjunctions, and product units with conjunctions. Graph \mathcal{G} can be consequently seen as a logical circuit encoding a Boolean formula. Both logical and probabilistic elements are considered by the PSDD definition provided in the next section.

2.2.2 Defining PSDDs

PSDDs have been originally presented as a probabilistic extension of a class of logical circuits called *sentential decision*

diagrams [Darwiche, 2011]. The extension is obtained by re-specifying the formula in the underlying circuit. Thus, unlike PCs, the position of the variables in a PSDD is constrained by a meta-structure called *vtree*.

A *vtree* over \mathbf{X} is just a full binary tree whose leaves are in one-to-one correspondence with the variables in \mathbf{X} (see, e.g., Figure 1). We intend *vtrees* as recursive objects and denote them with their root node. In practice, each node u of *vtree* v is a (sub-)vtree. Notation v is used for the variables of v . Moreover, we denote as v^l and v^r the sub-*vtrees* of v rooted at its left and right children.




Figure 1: Two *vtrees* over $\mathbf{X} = \{X_1, X_2, X_3, X_4\}$.

Using PCs terminology, we can define a PSDD associated to a *vtree* v by the properties of its structure and its parameters.

PSDD Structure. PSDD graph \mathcal{G} is defined as follows.

- Input units are associated with the leaves of *vtree* v , i.e., have scope $\{X\}$ for some $X \in v$. They can be literals (X or $\neg X$) or the logical constant \top (*true*).
- A product unit n has two inputs. Notation $n = (p, s)$ is used, p is called its *prime* and s its *sub*.
- A sum unit n is associated with an internal node u of v , its inputs are product units and are called its *elements*. The notation $n := \{(p_i, s_i)\}_{i=1}^k$ is used. Its primes and subs are either sum or input units. The scope of n and of its elements is u . The scope of the primes is u^l and the scope of the subs is u^r .
- The root is a sum unit.

PSDD Parameters. Parameters θ are defined as follows.

- Each input unit \top associated with variable X is annotated with a weight $\theta \in]0, 1[$.
- Each sum unit $\{(p_i, s_i)\}_{i=1}^k$ is annotated with non-negative weights $\theta_1, \dots, \theta_k$ such that $\sum_{i=1}^k \theta_i = 1$.

PSDD Logical Base Each PSDD node n encodes a Boolean formula $\langle n \rangle$ called its (logical) *base* and defined recursively as follows.

- For input units: $\langle X \rangle := X$, $\langle \neg X \rangle := \neg X$, and $\langle \top \rangle := \top$




Figure 2: A PSDD over four variables trained from the database in Table 1. Its *vtree* is in Figure 1a.

- If $n = (p, s)$ is a product unit, $\langle n \rangle := \langle p \rangle \wedge \langle s \rangle$
- If n is a sum unit, $n = \{n_i\}_{i=1}^k$, $\langle n \rangle := \bigvee_{i=1}^k \langle n_i \rangle$.

In fact, product units represent AND gates and sum units represent OR gates. The primes of a sum unit are defined as exclusive, i.e., $\langle p_i \rangle \wedge \langle p_j \rangle \equiv \perp$ for $i \neq j$.

Induced Probability Mass Function. Each unit n induces a probability mass function \mathbb{P}_n over the variables of its scope.

- If n is an input unit with scope $\{X\}$, \mathbb{P}_n is a univariate mass function: $\mathbb{P}_{\top}(X) := (\theta, 1 - \theta)$; $\mathbb{P}_X(X) := (1, 0)$; and $\mathbb{P}_{\neg X}(X) := (0, 1)$.
- Otherwise, let (\mathbf{X}, \mathbf{Y}) be the scope of n , where \mathbf{X} are its left variables and \mathbf{Y} its right ones.
 - If $n = (p, s)$ is a product unit,

$$\mathbb{P}_n(\mathbf{x}, \mathbf{y}) := \mathbb{P}_p(\mathbf{x})\mathbb{P}_s(\mathbf{y}). \quad (3)$$

- If $n = \{n_i\}_{i=1}^k$ is a sum unit,

$$\mathbb{P}_n(\mathbf{x}, \mathbf{y}) := \theta_i \mathbb{P}_{n_i}(\mathbf{x}, \mathbf{y}), \quad (4)$$

for the unique $1 \leq i \leq k$ such that $\mathbf{x} \models \langle p_i \rangle$. Moreover, each input of a sum unit assigns non-zero probability to at least one world.¹

¹This means that, unlike in the original definition, and as done in [Liang et al., 2017], we do not consider elements encoding the false, aka *dead branches*.

Let us note that the exclusivity of the primes makes PSDDs a special class of *deterministic* PCs. Indeed, for each decision node $n = \{(p_i, s_i)\}_{i=1}^k$ and for each world $(\mathbf{x}y)$, there is at most one prime p_i such that $\mathbb{P}_{p_i}(\mathbf{x}) > 0$. Moreover, the induced probability mass function of a PSDD factorizes through its variables by following its vtree. This makes PSDDs structured-decomposable PCs.

As a demonstrative example, Figure 2 depicts a PSDD over four variables consistent with the vtree in Figure 1.a and whose parameters have been trained from the database in Table 1 as done in [Kisa et al., 2014]. Labels on the decision nodes denote the vtree nodes for which the sub-PSDD is associated.

#	X_1	X_2	X_3	X_4
3	0	0	1	1
7	0	0	0	0
2	1	0	0	1
3	0	1	1	1
9	0	1	1	0
2	1	1	0	1
4	1	1	1	0
0	1	0	1	1
0	1	0	1	0
0	0	1	0	1

Table 1: A database for the joint states of four Boolean variables. The last three records are unobserved.

2.3 LEARNSPN

To conclude the review of the background material, let us consider the pseudo-code in Algorithm 1 where an adaptation of the classical LEARNSPN scheme for the learning of PCs from a database is depicted [Gens and Domingos, 2013]. This is a recursive procedure, where the database subject to horizontal (line 4) and vertical (line 6) splits. By horizontal split we intend a partition in clusters of similar records based on a clustering algorithm generically denoted as CLUSTER and returning the index sets of the different clusters. A vertical split is intended instead a partition of the variables in the database achieved by some independence test and denoted as PARTITION. When single columns are found (line 1), univariate mass functions are directly obtained from the data. The recursion is achieved by expressing the output mass function as a weighted sum of the mass functions associated to the different clusters (sums in line 8) and expressing them as a product of the mass functions associated with the different groups of variables (products in line 8).

Algorithm 1 LEARNSPN(\mathcal{D}, \mathbf{X}) (adaptated from Gens and Domingos [2013])

Return a PC given database \mathcal{D} over \mathbf{X}

```

1: if  $|\mathbf{X}| = 1$  then
2:   return  $\pi(\mathbf{X}) \leftarrow \mathcal{D}$ 
3: else
4:    $\{J_i\}_{i=1}^k \leftarrow \text{CLUSTER}(\mathcal{D}^{\mathbf{X}})$ 
5:   for  $i \leftarrow 1, \dots, k$  do
6:      $\{\mathbf{X}_i^j\}_{j=1}^{m_i} \leftarrow \text{PARTITION}(\mathbf{X}, \mathcal{D}_{J_i})$ 
7:   end for
8:   return  $\sum_{i=1}^k \frac{|\mathcal{D}_{J_i}|}{|\mathcal{D}|} \prod_{j=1}^{m_i} \text{LEARNSPN}(\mathcal{D}_{J_i}^{\mathbf{X}_i^j}, \mathbf{X}_i^j)$ 
9: end if

```

3 LEARNING PSDDS UNDER PCWA

The goal of this section is to show how the LEARNSPN scheme as sketched in Section 2.3 can be modified in order to force the output of Algorithm 1 to be a valid PSDD. Remarkably, this will be shown to correspond to a PCWA approach. Similarly to LEARNSPN, we might obtain a PSDD by recursively perform horizontal and vertical splits of \mathcal{D} . Yet, unlike LEARNSPN, the horizontal splits should be driven by the vtree, which we assume here available as an input.

Algorithm 2 SLOPP(\mathcal{D}, v)

Return a PSDD given database \mathcal{D} and vtree v and over \mathbf{X}

```

1: if  $|v| = 1$  then
2:    $X \leftarrow$  unique variable in  $v$ 
3:   if  $|\mathcal{D}^{X=1}| = |\mathcal{D}|$  then
4:     return  $X$ 
5:   else if  $|\mathcal{D}^{X=0}| = |\mathcal{D}|$  then
6:     return  $\neg X$ 
7:   else
8:     return  $(\top, \frac{|\mathcal{D}^{X=1}|}{|\mathcal{D}|})$ 
9:   end if
10: else
11:    $\{J_i\}_{i=1}^k \leftarrow \text{CLUSTER}(\mathcal{D}^{v^l})$ 
12:   for  $i = 1, \dots, k$  do
13:      $p_i \leftarrow \text{SLOPP}(\mathcal{D}_{J_i}^{v^l}, v^l)$ 
14:      $s_i \leftarrow \text{SLOPP}(\mathcal{D}_{J_i}^{v^r}, v^r)$ 
15:   end for
16:   return  $\{(p_i, s_i, \frac{|\mathcal{D}_{J_i}|}{|\mathcal{D}|})\}_{i=1}^k$ 
17: end if

```

We call our procedure SLOPP (structural learning of PSDDS under PCWA). Algorithm 2 depicts the SLOPP workflow. This is a recursive procedure, that starts growing the circuit from the top. A clustering algorithm (line 11) is executed on the columns of the database corresponding to the left variables of the vtree. The clusters obtained in this way are used to generate the primes (line 13), while the subs are made of the records for the right variables induced by

the same clusters (line 14). The probabilities assigned to these product units are proportional to the cardinality of the corresponding clusters (line 16). When we finally obtain databases over single variables (lines 2-8) univariate mass functions or simple literals are specified depending on the frequencies of the two Boolean states in the column. The only parameters of the algorithm are the number of clusters to be returned (or the criteria to select this number) and the minimum amount of records required to run the clustering algorithm. If the number of records is smaller than this threshold, a single cluster with all the records is returned.

As discussed in Section 2, given a database \mathcal{D} , CWA allows to describe the possible states of the affairs induced by the database as a DNF $\phi(\mathcal{D})$. The PSDD we learn represents a *relaxation* of $\phi(\mathcal{D})$, i.e., $\phi(\mathcal{D})$ logically implies the formula encoded by the learned PSDD. This is formalised by the following result, whose derivation is detailed in the appendix.

Proposition 1 *Algorithm 2 returns a valid PSDD representing a relaxation of the formula induced by the input database, i.e., $\phi(\mathcal{D})$ logically implies $\langle \text{SLOPP}(\mathcal{D}, v) \rangle$.*

The aforementioned relaxation happens because SLOPP applies a PCWA at each recursive call, as we shall explain with a simple example. Consider Table 1 and vtree v in Figure 1a whose left variables are $v^l = \{X_1, X_2\}$ and whose right variables are $v^r = \{X_3, X_4\}$. The horizontal lines in that table separate the output of a clustering of the records in \mathcal{D}^{v^l} , these inducing the primes of the root node, whose corresponding partition \mathcal{D}^{v^r} induces its subs. For instance, the second element is induced by the records corresponding to the third, fourth and fifth line in Table 1. In particular prime p_2 is induced by the sub-table described by the third, fourth and fifth lines restricted to the first two columns (associated to X_1 and X_2), whereas the restriction of those three lines to the last two columns (associated to X_3 and X_4) induces the corresponding sub s_2 . By combining the records of these two sub-tables, we see that three new *virtual records*, $[1, 0, 1, 1]$, $[1, 0, 1, 0]$ and $[0, 1, 0, 1]$ arise. Those three virtual records have zero counts, as showed at the bottom of Table 1. This notwithstanding, they are considered to be possible by the PSDD returned by SLOPP. Indeed, this PSDD, depicted in Figure 2, induces a mass function that assigns non-zero probability exactly to all the instances appearing in \mathcal{D} , included the virtual records. Stated otherwise, these records coincide with the joint states satisfying the Boolean formula encoded by the PSDD that is returned by SLOPP when applied to Table 1.

4 PRELIMINARY EXPERIMENTS

A first implementation of Algorithm 2 has been achieved within the Juice² (Julia) library for PCs [Dang et al., 2021]. The code, freely available as a simple Julia notebook,³ is highly experimental and not optimized for fast performance. For this reason, in this position paper, we discuss the results of a very preliminary validation based only on two small databases as available in Juice and whose features are detailed in Table 2. A deeper analysis based on a larger benchmark is a necessary future work.

Name	$ \mathcal{D}_{\text{train}} $	$ \mathcal{D}_{\text{test}} $	$ \mathbf{X} $
Nltcs	16181	3236	16
Plants	17412	3482	69

Table 2: Characteristics of the benchmark databases.

Given a complete database of Boolean variables, SLOPP needs a vtree over those variables. Here we adopt the state-of-the-art techniques proposed by Dang et al. [2021] and based on Chow-Liu trees as included in the (Juice implementation of the) STRUDEL algorithm. The PSDD returned by this algorithm is used to evaluate the performance of our method, together with a fully-factorised model intended to provide a trivial baseline level.

To cluster records we use the k-means algorithm. Here we set the number of clusters k constant (and equal to two and three), while a threshold on the database size d to create multiple clusters is also specified (and if this is not the case a single cluster with all the data is used).

As a consequence of Proposition 1, the PSDD returned by SLOPP might be consistent even with records not present in the training set. Yet, it might be possible that testing records not available in the training set would be inconsistent with the PSDD. We discard those records from the test database and we denote as γ their number. This clearly gives an unfair advantage to SLOPP if compared with STRUDEL or with the fully-factorised model, as both these PSDDs encode the tautology (i.e., no KB). Yet, for such a preliminary study, we are interested in evaluating the strength of such advantage in terms of performance or its relevance as shown by γ . Tables 3 and 4 depict the test log-likelihood on the consistent records for different values of the parameters.

As a positive remark we notice comparable performance levels between SLOPP and STRUDEL and a relatively low number of inconsistent test records. The most critical point seems to be the huge size of the PSDDs returned by SLOPP: for Plants there are more than an order of magnitude larger than the ones produced by STRUDEL. This is probably related to the fact that the current implementation of our

²github.com/Juice-jl.

³github.com/IDSIA-papers/2021-TPM.

\mathcal{C}	k	d	γ	$LL_{\mathcal{C}}$	$ \mathcal{C} $
FULLY FACT.	-	-	-	-35'844	79
STRUDEL	-	-	-	-20'054	786
SLOPP	2	20	4	-22'148	1'229
SLOPP	2	50	7	-23'441	1'232
SLOPP	3	20	23	-19'744	2'258
SLOPP	3	50	15	-20'176	2'033

Table 3: Test log-likelihood with Nltcs for SLOPP, STRUDEL and fully factorised models. Different experiments with different number of clusters k and threshold on the minimum cluster size d are reported. Parameter γ is the number of inconsistent test instances, while $|\mathcal{C}|$ is the number of nodes of the circuit \mathcal{C} obtained with the different approaches.

\mathcal{C}	k	d	γ	$LL_{\mathcal{C}}$	$ \mathcal{C} $
FULLY FACT.	-	-	-	-133'389	344
STRUDEL	-	-	-	-38'266	2'503
SLOPP	2	20	582	-40'442	71'602
SLOPP	2	50	594	-39'354	69'529
SLOPP	3	20	713	-48'329	103'742
SLOPP	3	50	793	-36'003	95'889

Table 4: Test log-likelihood with Plants. The same setup as in Table 3 is considered.

algorithm is only coping with singly-connected topologies, thus creating unnecessarily large circuits. Forcing the creation of multiple connections in the circuit could lead to smaller circuits.

5 CONCLUSIONS

A new structural learning algorithm for PSDDs has been presented. The algorithm uses the input data as a knowledge base to be relaxed by a PCWA approach. A very preliminary validation suggests that the algorithm might lead to reliable models to be used in frameworks where (some form of) PCWA is a tenable hypothesis. Besides an extensive experimental analysis, as a future work we intend to embed the learning of the vtree within the SLOPP architecture in order to better exploit the statistical information in the data set. In order to keep under control the size of the circuit returned by the algorithm, the dynamic generation of multiple connections when more elements involving the same sub-circuit are present should be also considered. A relaxed version of this idea might be used to obtain even smaller models. An integration between SLOPP and LEARNPSDD (e.g., achieved by using the output of the first algorithm as an input for the second) might also be considered to train better models. Finally, belief revision techniques could be considered to solve the training-testing-set inconsistency issue discussed in the experimental section.

PROOFS

Proof of Proposition 1 It is an easy exercise to verify that the output of SLOPP is a valid PSDD.

To see that the formula encoded by $\text{SLOPP}(\mathcal{D}, v)$ is a relaxation of $\phi(\mathcal{D})$, we proceed by induction on the number n of variables in \mathcal{D} . If $n = 1$, v is a leaf and $\text{SLOPP}(\mathcal{D}, v)$ is a trivial PSDD consisting in a single unit X , $\neg X$ or $(\top, \frac{|\mathcal{D}^{x=1}|}{|\mathcal{D}|})$. In each case the encoded formula is by definition $\phi(\mathcal{D})$. Assume that the statement holds for a number of variables strictly smaller than a fixed $n > 1$. Then, for a database \mathcal{D} over n variables and associated vtree v , consider $\text{SLOPP}(\mathcal{D}, v) = \{(p_i, s_i, \theta_i)\}_{i=1}^k$. For $1 \leq i \leq k$, $p_i = \text{SLOPP}(\mathcal{D}_{J_i}^l, v^l)$ and $s_i = \text{SLOPP}(\mathcal{D}_{J_i}^r, v^r)$, where J_1, \dots, J_k are the clusters obtained from \mathcal{D}^l .

Moreover, $\langle \text{SLOPP}(\mathcal{D}, v) \rangle = \bigvee_{i=1}^k \langle p_i \rangle \wedge \langle s_i \rangle$, and by induction hypothesis $\langle p_i \rangle$ and $\langle s_i \rangle$ are relaxations of $\phi(\mathcal{D}_{J_i}^l)$ and $\phi(\mathcal{D}_{J_i}^r)$, respectively.

This means that $\langle p_i \rangle$ implies $\phi(\mathcal{D}_{J_i}^l)$ and $\langle s_i \rangle$ implies $\phi(\mathcal{D}_{J_i}^r)$, and thus we get that $\langle p_i \rangle \wedge \langle s_i \rangle$ implies their conjunction. We can immediately conclude that $\langle \text{SLOPP}(\mathcal{D}, v) \rangle$ implies $\bigvee_{i=1}^k \phi(\mathcal{D}_{J_i}^l) \wedge \phi(\mathcal{D}_{J_i}^r)$.

Now, for $1 \leq i \leq k$, both $\phi(\mathcal{D}_{J_i}^l)$ and $\phi(\mathcal{D}_{J_i}^r)$ are DNFs whose conjunctive clauses are induced by the rows of $\mathcal{D}_{J_i}^l$ and of $\mathcal{D}_{J_i}^r$, respectively. By taking their conjunction we combine their clauses and clearly get, among others, $\phi(\mathbf{x}^i)$, the conjunctive clause induced by the i^{th} record of \mathcal{D} . This yields the desired conclusion.

References

- Franz Baader, Diego Calvanese, Deborah McGuinness, Peter Patel-Schneider, Daniele Nardi, et al. *The description logic handbook: Theory, implementation and applications*. Cambridge University Press, 2003.
- Yoo Jung Choi, Antonio Vergari, and Guy Van den Broeck. Probabilistic circuits: A unifying framework for tractable probabilistic models. Technical report, UCLA, 2020. URL <http://starai.cs.ucla.edu/papers/ProbCirc20.pdf>.
- Meihua Dang, Antonio Vergari, and Guy Broeck. Strudel: Learning structured-decomposable probabilistic circuits. In *International Conference on Probabilistic Graphical Models*, pages 137–148. PMLR, 2020.
- Meihua Dang, Pasha Khosravi, Yitao Liang, Antonio Vergari, and Guy Van den Broeck. Juice: A Julia package for logic and probabilistic circuits. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (Demo Track)*, Feb 2021.

- Fariz Darari, Werner Nutt, Giuseppe Pirro, and Simon Razniewski. Completeness statements about RDF data sources and their use for query answering. In *International Semantic Web Conference*, pages 66–83. Springer, 2013.
- Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 601–610, 2014.
- Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian Suchanek. AMIE: association rule mining under incomplete evidence in ontological knowledge bases. In *Proceedings of the 22nd international conference on World Wide Web*, pages 413–422, 2013.
- Robert Gens and Pedro Domingos. Learning the Structure of Sum-Product Networks. In *International conference on machine learning*, 2013.
- Tomasz Imieliński and Witold Lipski Jr. Incomplete information in relational databases. In *Readings in Artificial Intelligence and Databases*, pages 342–360. Elsevier, 1989.
- Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. Probabilistic sentential decision diagrams. In *Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2014.
- Yitao Liang, Jessa Bekker, and Guy Van den Broeck. Learning the Structure of Probabilistic Sentential Decision Diagrams. In *UAI*, 2017.
- Daniel Lowd and Pedro Domingos. Learning arithmetic circuits. *arXiv preprint arXiv:1206.3271*, 2012.
- Amihai Motro. Integrity = validity + completeness. *ACM Transactions on Database Systems (TODS)*, 14(4):480–502, 1989.
- Robert Peharz, Steven Lang, Antonio Vergari, Karl Stelzner, Alejandro Molina, Martin Trapp, Guy Van den Broeck, Kristian Kersting, and Zoubin Ghahramani. Einsum networks: Fast and scalable learning of tractable probabilistic circuits. In *International Conference on Machine Learning*, pages 7563–7574. PMLR, 2020.
- Hoifung Poon and Pedro Domingos. Sum-Product Networks: A New Deep Architecture. In *ICCV*. IEEE, 2011.
- Raymond Reiter. On closed world data bases. In *Readings in artificial intelligence*, pages 119–140. Elsevier, 1981.