

Probabilistic Sentential Decision Diagrams

Doga Kisa and Guy Van den Broeck and Arthur Choi and Adnan Darwiche

Computer Science Department

University of California, Los Angeles

{doga, guyvdb, aychoi, darwiche}@cs.ucla.edu

Abstract

We propose the Probabilistic Sentential Decision Diagram (PSDD): A complete and canonical representation of probability distributions defined over the models of a given propositional theory. Each parameter of a PSDD can be viewed as the (conditional) probability of making a decision in a corresponding Sentential Decision Diagram (SDD). The SDD itself is a recently proposed complete and canonical representation of propositional theories. We explore a number of interesting properties of PSDDs, including the independencies that underlie them. We show that the PSDD is a tractable representation. We further show how the parameters of a PSDD can be efficiently estimated, in closed form, from complete data. We empirically evaluate the quality of PSDDs learned from data, when we have knowledge, a priori, of the domain logical constraints.

Introduction

The interplay between logic and probability has been of great interest throughout the history of AI. One of the earliest proposals in this direction is Nilsson’s (1986) probabilistic logic, which aimed at augmenting first-order logic with probabilities. This has prompted similar approaches, including, for example, Halpern (1990). The focus of these approaches, however, was mainly semantical, yielding no effective schemes for realizing them computationally. More recently, the area of lifted probabilistic inference has tackled this interplay, while employing a different compromise (Poole 2003). In these efforts, the focus has been mostly on restricted forms of first-order logic (e.g., function-free and finite domain), but with the added advantage of efficient inference (e.g., algorithms whose complexity is polynomial in the domain size).

On the propositional side, the thrust of the interplay has been largely computational. An influential development in this direction has been the realization that enforcing certain properties on propositional representations, such as decomposability and determinism, provides one with the power to answer probabilistic queries efficiently. This development was actually based on two technical observations. First, that decomposable and deterministic representations allow one to perform weighted model counting efficiently. Second, that probabilistic reasoning can be reduced to weighted model

counting. This development, which has its first roots in Darwiche (2002), has been underlying an increasing number of probabilistic reasoning systems in the last decade. This is especially true for representations that employ both logical and probabilistic elements (e.g., Chavira, Darwiche, and Jaeger (2006) and Fierens et al. (2011)). Moreover, the technique has been extended recently to certain first-order representations as well (Van den Broeck et al. 2011).

This paper is concerned with an orthogonal contribution to this interplay between propositional logic and probability theory. The problem we tackle here is that of developing a representation of probability distributions in the presence of massive, logical constraints. That is, given a propositional logic theory which represents domain constraints, our goal is to develop a representation that induces a unique probability distribution over the models of the given theory. Moreover, the proposed representation should satisfy requirements that are sometimes viewed as necessary for the practical employment of such representations. These include a clear semantics of the representation parameters; an ability to reason with the representation efficiently; and an ability to learn its parameters from data, also efficiently.

Our proposal is called a *Probabilistic Sentential Decision Diagram (PSDD)*. It is based on the recently proposed *Sentential Decision Diagram (SDD)* for representing propositional theories (Darwiche 2011; Xue, Choi, and Darwiche 2012; Choi and Darwiche 2013). While the SDD is comprised of *logical* decision nodes, the PSDD is comprised of *probabilistic* decision nodes, which are induced by supplying a distribution over the branches of a logical decision node. Similar to SDDs, the PSDD is a canonical representation, but under somewhat more interesting conditions. Moreover, computing the probability of a term can be done in time linear in the PSDD size. In fact, the probability of each and every literal can be computed in only two passes over the PSDD. It is particularly notable that the local parameters of a PSDD have clear semantics with respect to the global distribution induced by the PSDD. We will also show that these parameters can be learned efficiently from complete data.

This paper is structured as follows. We start by a concrete discussion on some of the applications that have driven the development of PSDDs and follow by an intuitive exposure of PSDDs and their salient features. We next provide a formal treatment of the syntax, semantics and properties of

L	K	P	A	Students
0	0	1	0	6
0	0	1	1	54
0	1	1	1	10
1	0	0	0	5
1	0	1	0	1
1	0	1	1	0
1	1	0	0	13
1	1	1	0	8
1	1	1	1	3

Table 1: Student enrollment data.

PSDDs. This allows us to present the main inference algorithm for PSDDs and the one for learning PSDD parameters from complete data. The paper concludes with some experimental results showing the promise of PSDDs in learning probability distributions under logical constraints. Proofs of theorems are delegated to the full version of the paper due to space limitations.

Motivation

PSDDs were inspired by the need to learn probability distributions that are subject to domain constraints. Take for example a computer science department that organizes four courses: Logic (L), Knowledge Representation (K), Probability (P), and Artificial Intelligence (A). Students are asked to enroll for these courses under the following restrictions:

- A student must take at least one of Probability or Logic.
- Probability is a prerequisite for AI.
- The prerequisite for KR is either AI or Logic.

The department may have data on student enrollments, as in Table 1, and may wish to learn a probabilistic model for reasoning about student preferences. For example, the department may need to know whether students are more likely to satisfy the prerequisite of KR using AI or using Logic.

A mainstream approach for addressing this problem is to learn a probabilistic graphical model, such as a Bayesian network. In this case, a network structure is constructed manually or learned from data. The structure is then turned into a Bayesian network by learning its parameters from the data. Other graphical models can also be used. This includes, for example, Markov networks or their variations.

What is common among all these approaches is that they lack a principled and effective method for accommodating the domain constraints into the learning process—that is, ensuring, for example, that a student with a profile $A \wedge K \wedge L \wedge \neg P$, or a profile $\neg A \wedge K \wedge \neg L \wedge P$, has zero probability in the learned model. In principle, the zero parameters of a graphical model can capture logical constraints, although a fixed model structure will not in general accommodate arbitrary logical constraints. We could introduce additional structure into the model to capture such constraints, using, e.g., the method of virtual evidence (Pearl 1988; Mateescu and Dechter 2008). However, incorporating constraints in this manner will in general lead to a highly-connected network, making inference intractable. Even if inference remained tractable, such an approach is not ideal as

we now have to learn a distribution that is conditioned on the constraints. This would require new learning algorithms (e.g., gradient methods) for performing parameter estimation as traditional methods may no longer be applicable. For example, in Bayesian networks, the closed-form parameter estimation algorithm under complete data will no longer be valid in this case.

The domain constraints of our example can be expressed using the following CNF.

$$\begin{aligned} P \vee L \\ A \Rightarrow P \\ K \Rightarrow A \vee L \end{aligned} \tag{1}$$

Even though there are 16 combinations of courses, the CNF says that only 9 of them are valid choices. An approach that observes this information must learn a probability distribution that assigns a zero probability to every combination that is not allowed by these constraints.

None of the standard learning approaches we are familiar with has been posed to address this problem. The complication here is not strictly with the learning approaches, but with the probabilistic models that are amenable to being learned under these circumstances. In particular, these models are not meant to induce probability distributions that respect a given set of logical constraints.

The simple problem we posed in this section is exemplary of many real-world applications. We mention in particular *configuration* problems that arise when purchasing products, such as cars and computers. These applications give users the option to configure products, but subject to certain constraints. Data is abundant for these applications and there is a clear economic interest in learning probabilistic models under the given constraints. We also mention *reasoning about physical systems*, which includes verification and diagnosis applications. Here, propositional logic is typically used to encode some system functionality, while leaving out some system behaviors which may have a non-deterministic nature (e.g., component failures and probabilistic transitions). There is also an interest here to learn probabilistic models of these systems, subject to the given constraints.

Our goal in this paper is to introduce the PSDD representation for addressing this particular need. We will start by an intuitive (and somewhat informal) introduction to PSDDs, followed by a more formal treatment of their syntax, semantics and the associated reasoning and learning algorithms.

PSDDs

We will refer to domain constraints as the *base* of a probability distribution. Our proposed approach starts by representing this base as a *Sentential Decision Diagram (SDD)* as in Figure 1 (Darwiche 2011; Xue, Choi, and Darwiche 2012; Choi and Darwiche 2013). An SDD is determined by a *vtree*, which is a full binary tree with leaves corresponding to the domain variables (Pipatsrisawat and Darwiche 2008). The choice of a particular SDD can then be thought of as a choice of a particular vtree. We will later discuss the impact of this choice on the represented distribution. For now, however, we will develop some further understanding of SDDs as they are the backbones of our probability distributions.

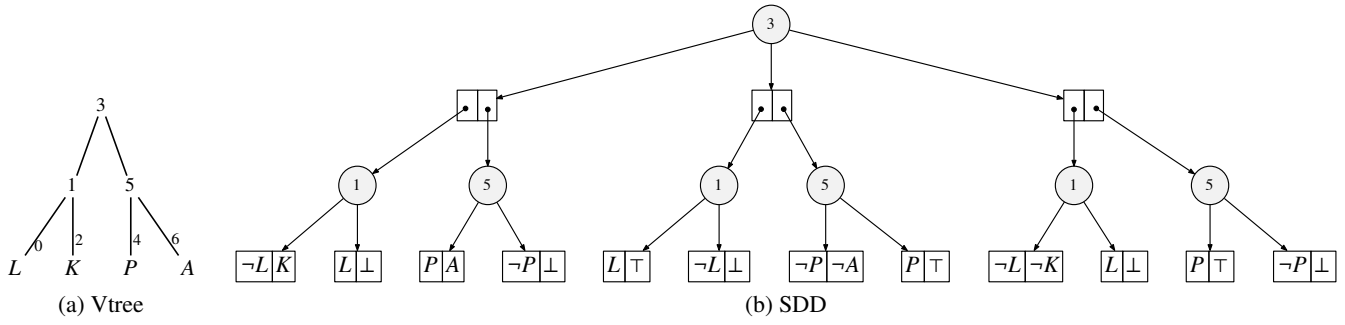


Figure 1: A vtree and SDD for the student enrollment problem. Numbers in circles correspond to vtree nodes.

L	K	P	A	Students	Learned PSDD Distribution
0	0	1	0	6	$0.6 \cdot 0.1$ 6.0%
0	0	1	1	54	$0.6 \cdot 0.9$ 54.0%
0	1	1	1	10	0.1 10.0%
1	0	0	0	5	$0.3 \cdot 0.2 \cdot 0.6$ 3.6%
1	0	1	0	1	$0.3 \cdot 0.2 \cdot 0.4 \cdot 0.75$ 1.8%
1	0	1	1	0	$0.3 \cdot 0.2 \cdot 0.4 \cdot 0.25$ 0.6%
1	1	0	0	13	$0.3 \cdot 0.8 \cdot 0.6$ 14.4%
1	1	1	0	8	$0.3 \cdot 0.8 \cdot 0.4 \cdot 0.75$ 7.2%
1	1	1	1	3	$0.3 \cdot 0.8 \cdot 0.4 \cdot 0.25$ 2.4%

Table 2: Student enrollment data and learned distribution.

SDDs. An SDD is either a *decision node* or a *terminal node*. A terminal node is a literal, the constant \top (true) or the constant \perp (false). A decision node is a disjunction of the form $(p_1 \wedge s_1) \vee \dots \vee (p_n \wedge s_n)$, where each pair (p_i, s_i) is called an *element*. A decision node is depicted by a circle and its elements are depicted by paired boxes. Here, p_1, \dots, p_n are called *primes* and s_1, \dots, s_n are called *subs*. Primes and subs are themselves SDDs. Moreover, the primes of a decision node are always consistent, mutually exclusive and exhaustive. The SDD in Figure 1 has seven decision nodes. The decision node to the far left has two elements $(\neg L, K)$ and (L, \perp) . It represents $(\neg L \wedge K) \vee (L \wedge \perp)$, which is equivalent to $\neg L \wedge K$. There are two primes for this node $\neg L$ and L . The two corresponding subs are K and \perp .

Structure. An SDD can be viewed as a *structure* that induces infinitely many probability distributions (all having the same base). By *parameterizing* an SDD, one obtains a PSDD that induces a particular probability distribution.

Parameters. Figure 2 depicts a PSDD which is obtained by parameterizing the SDD in Figure 1. Both decision and terminal SDD nodes are parameterized, but we focus here on decision nodes. Let n be a decision node having elements $(p_1, s_1), \dots, (p_n, s_n)$. To parameterize node n is to provide a distribution $\theta_1, \dots, \theta_n$. Intuitively, θ_i is the probability of prime p_i given that the decision of node n has been implied. We will formalize and prove this semantics later. We will also provide an efficient procedure for learning the parameters of a PSDD from complete data. The PSDD parameters in Figure 2 were learned using this procedure from the data in Table 2. The table also depicts the probability distribution induced by the learned PSDD.

Independence. The PSDD structure is analogous to a Bayesian network structure in the following sense. The latter can be parameterized in infinitely many ways, with each parameterization inducing a particular probability distribution. Moreover, all the induced distributions satisfy certain independences that can be inferred from the underlying Bayesian network structure. The same is true for PSDDs. Each parameterization of a PSDD structure yields a unique probability distribution. Moreover, all the induced distributions satisfy independences that can be inferred from the PSDD structure. We will show, however, that PSDD independence is more refined than Bayesian network independence as it allows one to express more qualified independence statements.

The Syntax and Semantics of PSDDs

PSDDs are based on *normalized SDDs* in which every node n is associated with (normalized for) a vtree node v according to the following rules (Darwiche 2011).

- If n is a terminal node, then v is a leaf node which contains the variable of n (if any).
- If n is a decision node, then its primes (subs) are normalized for the left (right) child of v .
- If n is the root SDD node, then v is the root vtree node.

The SDD in Figure 1 is normalized. Each decision node in this SDD is labeled with the vtree node it is normalized for. We are now ready to define the syntax of a PSDD.

Definition 1 (PSDD Syntax) A PSDD is a normalized SDD with the following parameters.

- For each decision node $(p_1, s_1), \dots, (p_k, s_k)$ and prime p_i , a positive parameter θ_i is supplied such that $\theta_1 + \dots + \theta_k = 1$ and $\theta_i = 0$ iff $s_i = \perp$.
- For each terminal node \top , a positive parameter θ is supplied such that $0 < \theta < 1$.

A terminal node \top with parameter θ will be denoted by $X:\theta$, where X is the variable of leaf vtree node that \top is normalized for. Other terminal nodes (i.e., \perp , X and $\neg X$) have fixed, implicit parameters (discussed later) and will be denoted as is. A decision PSDD node will be denoted by $(p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$. Graphically, we will just annotate the edge into element (p_i, s_i) with the parameter θ_i . Figure 2 provides examples of this notation.




Figure 2: A PSDD for the student enrollment problem, which results from parameterizing the SDD in Figure 1. The parameters were learned from the dataset in Table 1 (also shown in Table 2).

We next define the distribution of a PSDD, inductively. That is, we first define the distribution induced by a terminal node. We then define the distribution of a decision node in terms of the distributions induced by its primes and subs.

Definition 2 (PSDD Distribution) Let n be a PSDD node that is normalized for vtree node v . Node n defines a distribution \Pr_n over the variables of vtree v as follows.

– If n is a terminal node, and v has variable X , then

n	$\Pr_n(X)$	$\Pr_n(\neg X)$
$X:\theta$	θ	$1 - \theta$
\perp	0	0
X	1	0
$\neg X$	0	1

– If n is a decision node $(p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$ and v has left variables \mathbf{X} and right variables \mathbf{Y} , then

$$\Pr_n(\mathbf{xy}) \stackrel{\text{def}}{=} \Pr_{p_i}(\mathbf{x}) \cdot \Pr_{s_i}(\mathbf{y}) \cdot \theta_i \quad \text{for } i \text{ where } \mathbf{x} \models p_i.$$

Applying this definition to the PSDD of Figure 2 leads to the distribution in Table 2 for its root node. The following table depicts the distribution induced by a non-root node in this PSDD, which appears in the middle of Figure 2.

\mathbf{x}	\mathbf{y}	$\Pr_{p_i}(\mathbf{x})$	$\Pr_{s_i}(\mathbf{y})$	θ_i	$\Pr_n(\mathbf{xy})$
P	A	1	0.25	0.4	0.1
P	$\neg A$	1	0.75	0.4	0.3
$\neg P$	A	1	0	0.6	0.0
$\neg P$	$\neg A$	1	1	0.6	0.6

Table 3: Distribution of node $n = (\neg P, \neg A)(P, \top)$.

The SDD node associated with a PSDD node n is called the *base* of n and is denoted by $[n]$. When there is no ambiguity, we will often not distinguish between a PSDD node n and its base $[n]$.

A PSDD assigns a strictly positive probability to a variable instantiation iff the instantiation satisfies its base. This can be seen, for example, in Table 3. This is also the first key property of PSDDs.

Theorem 1 (Base) Consider a PSDD node n that is normalized for vtree node v . If \mathbf{Z} are the variables of vtree v , then $\Pr_n(\mathbf{z}) > 0$ iff $\mathbf{z} \models [n]$.

We will now discuss the second key property of PSDDs, which reveals the *local semantics* of PSDD parameters.

Theorem 2 (Parameter Semantics) Let n be a decision node $(p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$. We have $\theta_i = \Pr_n([p_i])$.

Consider the PSDD in Figure 2 and its decision node n in Table 3. Prime $\neg P$ of this node has parameter 0.6. According to Theorem 2, we must then have $\Pr_n(\neg P) = 0.6$, which can be verified in Table 3. Similarly, $\Pr_n(P) = 0.4$.

The third key property of PSDDs is the relationship between the local distributions induced by its various nodes (node distributions) and the global distribution induced by its root node (PSDD distribution)—for example, the relationship between the distribution of node n in Table 3 and the PSDD distribution given in Table 2.

Node distributions are linked to the PSDD distribution by the notion of *context*.

Definition 3 (Context) Let $(p_1, s_1), \dots, (p_k, s_k)$ be the elements appearing on some path from the SDD root to node n .¹ Then $p_1 \wedge \dots \wedge p_k$ is called a *sub-context* for node n and is feasible iff $s_i \neq \perp$. The *context* is a disjunction of all sub-contexts and is feasible iff some sub-context is feasible.

Consider Figure 1. The three decision nodes normalized for vtree node $v = 5$ have the contexts $\neg L \wedge K$, L and $\neg L \wedge \neg K$. Moreover, the terminal nodes normalized for vtree $v = 6$ have the contexts:

- A : $\neg L \wedge K \wedge P$
- $\neg A$: $L \wedge \neg P$
- \perp : $(\neg L \wedge K \wedge \neg P) \vee (\neg L \wedge \neg K \wedge \neg P) = (\neg L \wedge \neg P)$
- \top : $(L \wedge P) \vee (\neg L \wedge \neg K \wedge P) = (L \vee \neg K) \wedge P$.

Contexts satisfy interesting properties.

Theorem 3 (Context) A node is implied by its context and the underlying SDD. Nodes normalized for the same vtree node have mutually exclusive and exhaustive contexts. The

¹That is, $n = p_k$ or $n = s_k$.




Figure 3: A Bayesian network structure.

sub-contexts of a node are mutually exclusive. A context/sub-context is feasible iff it has a strictly positive probability.

Contexts give a global interpretation to node distributions.

Theorem 4 (Node Distribution) Consider a PSDD r and let n be one of its nodes. If γ_n is a feasible sub-context or feasible context of node n , then $\Pr_r(\cdot) = \Pr_r(\cdot | \gamma_n)$.

Contexts also give a global interpretation to parameters.

Corollary 1 (Parameter Semantics) Consider a PSDD r and node n with feasible sub-context or feasible context γ_n .

- If n is a terminal node $X : \theta$, then $\theta = \Pr_r(X | \gamma_n)$.
- If n is a decision node $(p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$, then $\theta_i = \Pr_r([p_i] | \gamma_n)$ for $i = 1, \dots, k$.

This corollary says that the parameters of a node are conditional probabilities of the PSDD distribution.

We show in the Appendix that PSDDs are *complete* as they are capable of representing any probability distribution. We also show that PSDDs are *canonical* under a condition known as compression. More precisely, we show that there is a unique compressed PSDD for each distribution and vtree. This is particularly important for learning PSDDs (structure and parameters) as it reduces the problem of searching for a PSDD into the problem of searching for a vtree.

PSDD Independence

Consider the Bayesian network structure in Figure 3, which corresponds to our earlier example. This structure encodes a number of probabilistic independences that hold in any distribution it induces (i.e., regardless of its parameters). These independences are

- A and L are independent given P .
- K and P are independent given AL .

These independences are conditioned on variables. That is, “given AL ” reads “given any state of variables A and L .” The second independence can therefore be expanded into $2 \times 2 \times 4$ statements of the form “ α is independent of β given γ ,” where α, β and γ are propositional sentences (e.g., $\neg K$ is independent of P given $A \wedge \neg L$).

The structure of a PSDD also encodes independences that hold in every induced distribution. These independences fall into two major categories, the first coming from Theorem 4.

Corollary 2 (Independence I) Consider a PSDD r and a node n with context γ_n and feasible sub-context β_n .

- If n is a terminal node $X : \theta$, then

$$\Pr_r(X | \gamma_n, \beta_n) = \Pr_r(X | \gamma_n) = \Pr_r(X | \beta_n) = \theta.$$

- If n is a decision node $(p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$, then

$$\Pr_r([p_i] | \gamma_n, \beta_n) = \Pr_r([p_i] | \gamma_n) = \Pr_r([p_i] | \beta_n) = \theta_i$$

for $i = 1, \dots, k$.

That is, the probability of a prime is independent of a sub-context once the context is known. This is also true for the probability of a terminal sub. Moreover, which specific sub-context we know is irrelevant. All are equivalent as far as defining the semantics of parameters is concerned.

The second category of independences is as follows.

Theorem 5 (Independence II) Let γ_v be a feasible context for a PSDD node normalized for vtree node v . Variables inside v are independent of those outside v given context γ_v .

To read the independences characterized by this theorem, one iterates over each vtree node v , identifying its corresponding, feasible contexts γ_v . Consider the PSDD in Figure 2. There are three decision nodes which are normalized for vtree node $v = 5$, with contexts $\neg L \wedge K$, L and $\neg L \wedge \neg K$. Using the second context, we get

variables PA and LK are independent given context L .

This reads as “If we know that someone took Logic, then whether they took KR has no bearing on whether they took Probability or AI.”

PSDD independences are conditioned on propositional sentences (contexts) instead of variables. This kind of independence is known to be more expressive and is usually called *context-specific independence* (Boutilier et al. 1996). This kind of independence is beyond the scope of probabilistic graphical models, which can only condition independence statements on variables. Consider the example statement we discussed above. If we were to condition on the variable L instead of the propositional sentence L , we would also get “If we know that someone *did not take* Logic, then whether they took KR has no bearing on whether they took Probability or AI.” This is actually contradicted by the logical constraints for this problem ($K \Rightarrow A \vee L$). If someone did not take Logic, but took KR, they must have taken AI.

Reasoning with PSDDs

We now present the main algorithms for reasoning with PSDDs. In particular, given a PSDD r and an instantiation e of some variables (*evidence*), we provide an algorithm for computing the probability of this evidence $\Pr_r(e)$. We also present an algorithm for computing the conditional probability $\Pr_r(X | e)$ for every variable X . Both algorithms run in time which is linear in the PSDD size.

We start with the first algorithm. For variable instantiation e and vtree node v , we will use e_v to denote the subset of instantiation e that pertains to the variables of vtree v , and $e_{\bar{v}}$ to denote the subset of e that pertains to variables outside v . The first algorithm is based on the following result.

Algorithm 1: Probability of Evidence

Input: PSDD r and evidence e

- 1 $evd(n) \leftarrow 0$ for every node n
// visit children before parents
- 2 **foreach** node n in the PSDD **do**
- 3 **if** n is a terminal node **then**
- 4 $v \leftarrow$ leaf vtree node that n is normalized for
- 5 $evd(n) \leftarrow \Pr_n(\mathbf{e}_v)$
- 6 **else**
- 7 **foreach** element (p_i, s_i, θ_i) of node n **do**
- 8 $evd(n) \leftarrow evd(n) + evd(p_i) \cdot evd(s_i) \cdot \theta_i$

Theorem 6 Consider a decision node $n = (p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$ that is normalized for vtree node v , with left child l and right child r . For evidence e , we have

$$\Pr_n(\mathbf{e}_v) = \sum_{i=1}^k \Pr_{p_i}(\mathbf{e}_l) \cdot \Pr_{s_i}(\mathbf{e}_r) \cdot \theta_i$$

When n is a terminal node, v is a leaf vtree node and \mathbf{e}_v is either a literal or the empty instantiation. In this case, we can just look up the value of $\Pr_n(\mathbf{e}_v)$ from the distribution induced by the terminal node n (from Definition 2).

Theorem 6 leads to Algorithm 1, which traverses the PSDD bottom-up, computing $\Pr_n(\mathbf{e}_v)$ for each node n and storing the result in $evd(n)$. The probability of evidence is then $evd(r)$, where r is the PSDD root.

We now turn to computing the probability $\Pr_r(X, e)$ for each variable X . One can use Algorithm 1 to perform this computation, but the algorithm would need to be called once for each variable X . However, with the following theorem, we can compute all of these node marginals using a single, second pass on the PSDD, assuming that Algorithm 1 did the first pass.

Theorem 7 Consider a PSDD r , variable X , and its leaf vtree node v . Let n_1, \dots, n_k be all the terminal nodes normalized for v and let $\gamma_{n_1}, \dots, \gamma_{n_k}$ be their corresponding contexts. For evidence e , we have

$$\Pr_r(X, \mathbf{e}_{\bar{v}}) = \sum_{i=1}^k \Pr_{n_i}(X) \cdot \Pr_r(\gamma_{n_i}, \mathbf{e}_{\bar{v}}).$$

If $e \models \neg X$, then $\Pr_r(X, e) = 0$. Otherwise, $X, \mathbf{e}_{\bar{v}} = X, e$ and $\Pr_r(X, e) = \Pr_r(X, \mathbf{e}_{\bar{v}})$.

The term $\Pr_{n_i}(X)$ in Theorem 7 is immediately available. Algorithm 2 computes $\Pr_r(\gamma_n, \mathbf{e}_{\bar{v}})$ for every PSDD node n that has context γ_n and is normalized for vtree node v . The algorithm traverses the PSDD top-down, computing this probability for each visited node n , storing it in $ctx(n)$. If n is a terminal node, the algorithm also computes $\Pr_r(X, \mathbf{e}_{\bar{v}})$ and $\Pr_r(\neg X, \mathbf{e}_{\bar{v}})$, storing them in $mrg(X)$ and $mrg(\neg X)$.

The simplicity of Algorithm 2 is due to the following. The probability of a sub-context can be computed by multiplying the parameters appearing on its corresponding path. Since sub-contexts are mutually exclusive, their probabilities can

Algorithm 2: Probability of Contexts

Input: PSDD r

- 1 $ctx(n) \leftarrow 0$ for nodes $n \neq r$ and $ctx(r) \leftarrow 1$
- 2 $mrg(X) \leftarrow 0$ and $mrg(\neg X) \leftarrow 0$ for every variable X
// visit parents before children
- 3 **foreach** node n in the PSDD **do**
- 4 **if** n is a terminal node **then**
- 5 $X \leftarrow$ variable of node n
- 6 $mrg(X) \leftarrow mrg(X) + ctx(n) \cdot \Pr_n(X)$
- 7 $mrg(\neg X) \leftarrow mrg(\neg X) + ctx(n) \cdot \Pr_n(\neg X)$
- 8 **else**
- 9 **foreach** element (p_i, s_i, θ_i) of node n **do**
- 10 $ctx(p_i) \leftarrow ctx(p_i) + \theta_i \cdot evd(s_i) \cdot ctx(n)$
- 11 $ctx(s_i) \leftarrow ctx(s_i) + \theta_i \cdot evd(p_i) \cdot ctx(n)$

be added to obtain the context probability. Algorithm 2 does precisely this except that it accounts for evidence as well using quantities computed by Algorithm 1.

Learning with PSDDs

We now present an algorithm for learning the parameters of a PSDD from a *complete* dataset. We start first with some basic definitions. An instantiation of all variables is called an *example*. There are 2^n distinct examples over n propositional variables. A complete dataset is a multi-set of examples.² That is, an example may appear multiple times in a dataset. Given a PSDD structure (a normalized SDD), and a complete dataset, our goal is to learn the value of each PSDD parameter. More precisely, we wish to learn *maximum likelihood* parameters: ones that maximize the probability of examples in the dataset.

We will use \Pr_θ to denote the distribution induced by the PSDD structure and parameters θ . The *likelihood* of these parameters given dataset \mathbf{D} is defined as

$$L(\theta|\mathbf{D}) = \prod_{i=1}^N \Pr_\theta(\mathbf{d}_i),$$

where \mathbf{d}_i ranges over all N examples in dataset \mathbf{D} . Our goal is then to find the maximum likelihood parameters

$$\theta^{ml} = \operatorname{argmax}_\theta L(\theta|\mathbf{D}).$$

We will use $\mathbf{D}\#(\alpha)$ to denote the number of examples in dataset \mathbf{D} that satisfy propositional sentence α . For a decision node $n = (p_1, s_1, \theta_1), \dots, (p_k, s_k, \theta_k)$ with context γ_n , we propose the following estimate for parameter θ_i :

$$\theta_i^{ml} = \frac{\mathbf{D}\#(p_i, \gamma_n)}{\mathbf{D}\#(\gamma_n)}. \quad (2)$$

For terminal node $n = X:\theta$ with context γ_n , we propose the following estimate for parameter θ :

$$\theta^{ml} = \frac{\mathbf{D}\#(X, \gamma_n)}{\mathbf{D}\#(\gamma_n)}. \quad (3)$$

²In an *incomplete* dataset, an example corresponds to an instantiation of some variables (not necessarily all).

We can now show the following.

Theorem 8 *The parameter estimates of Equations 2 and 3 are the only estimates that maximize the likelihood function.*

Our parameter estimates admit a *closed-form*, in terms of the counts $\mathbf{D}\#(\alpha)$ in the data. One can compute these estimates using a single pass through the examples of a dataset. Moreover, each distinct example can be processed in time linear in the PSDD size.³ These are very desirable properties for a parameter learning algorithm. These properties are shared with Bayesian network representations, but are missing from many others, including Markov networks.

When learning probabilistic graphical models, one makes a key distinction between learning structures versus learning parameters (the former being harder in general). While learning PSDD structures is beyond the scope of this paper, the experimental results we present next do use a basic method for learning structures. In particular, since we compile logical constraints into an SDD (i.e., a PSDD structure), the compilation technique we use is effectively “learning” a structure. We used the publicly available SDD package for this purpose (<http://reasoning.cs.ucla.edu/sddl/>). The SDD package tries to dynamically minimize the size of the compiled SDD and, as a result, tries to minimize the number of PSDD parameters.

Preliminary Experimental Results

In this Section, we empirically evaluate our parameter estimation algorithm for PSDDs. First, we illustrate how our algorithm can effectively recover the parameters of a PSDD from data, assuming that the dataset was indeed generated by the distribution of a PSDD. Second, we highlight how knowing the logical constraints underlying a given dataset can impact the accuracy of learned models.

In our first set of experiments, using synthetic data, we show how we can recover the true parameters that generated a dataset, given that the logical constraints are known. As we will be simulating datasets from a known PSDD, we shall use the KL-divergence to compare the original PSDD distribution $\Pr(\mathbf{Z})$, and the PSDD distribution $\Pr'(\mathbf{Z})$ that we learned from the data:

$$KL(\Pr, \Pr') = \sum_{\mathbf{z}} \Pr(\mathbf{z}) \log \frac{\Pr(\mathbf{z})}{\Pr'(\mathbf{z})}.$$

Note that the KL-divergence is non-negative, and zero iff the two distributions are equivalent.⁴

³A dataset may not include every example that is consistent with the domain constraints. If this is the case, a parameter θ for prime p may be estimated to zero, even though its sub s may not be \perp ; see Theorem 1. To address this, one can assume a pseudo-count for each distinct example, which can be thought of as providing a prior distribution on parameters. In our experiments, we assumed a pseudo-count of $1/mc$ for each distinct example, where mc is the model count of the SDD. This corresponds to a very weak prior since, in aggregate, these pseudo-counts contribute a total count that is equivalent to one real example in the dataset.

⁴The KL-divergence between two PSDDs can be computed efficiently if they share the same structure.




Figure 4: We observe how our parameter estimation algorithm can recover the original PSDD as we increase the size of the training dataset (x -axis), measured by the KL-divergence (y -axis, where lower is better).

Consider Figure 4, where each point in the plot represents an average of 100 randomly generated PSDDs. We generated 10 random PSDD structures by compiling random 3-CNFs into SDDs (the used clause to variable ratio is 3). For each SDD, we randomly generated 10 different parameterizations, giving us $10 \times 10 = 100$ PSDDs total. The SDDs were compiled using dynamic minimization, which reduces the size of the SDD, and hence, the number of parameters in the corresponding PSDD. We evaluated PSDDs over 16 variables and 32 variables (one curve each).

As we go right on the x -axis, we increase the size of the dataset that we are learning our PSDD from. We clearly see that as more data becomes available, the parameters that we learn more closely match the true parameters of the original PSDD (measured using the KL-divergence, on the y -axis). Indeed, both curves are steadily tending to zero. The PSDDs over 16 variables converge faster than those over 32 variables. The PSDDs over 16 variables have fewer parameters, roughly, a tenth as many.

Our final experiment tries to put PSDD learning in the context of learning probabilistic graphical models. It is hard to make a direct comparison here since existing approaches do not factor domain constraints into the learning process. Hence, there are no available systems that tackle the learning problem we posed in this paper, nor are there reported experimental results that are directly relevant. Still, we conceived an experiment that is revealing in this regard, which we explain next.

A standard experiment for evaluating a learning approach goes as follows. The dataset is divided into a *training* portion and *testing* one. The learning algorithm is then given the training dataset and asked to find a probabilistic model. The quality of the model is then measured by computing the likelihood of learned parameters, but using the testing dataset. The literature contains many experiments of this form (e.g., Table 4). The question we posed is this: What if the dataset embedded the domain constraints, in that any example that

Table 4: Log Likelihoods of Testing Datasets. A question mark indicates the unavailability of likelihood for the corresponding approach/dataset.

benchmark	#var	ACMN	ACBN	LTM	CLT	PSDD
NLTCS	16	-6.01	-6.02	-6.49	-6.76	-6.31
MSNBC	17	-6.04	-6.04	-6.52	-6.54	-6.20
KDDCup	64	-2.15	-2.16	-2.18	-2.29	-2.11
Plants	69	-12.89	-12.85	-16.39	-16.52	-10.99
Audio	100	-40.32	-41.13	-41.90	-44.37	-19.34
Jester	100	-53.35	-54.43	-55.17	-58.23	-18.49
Netflix	100	-57.26	-57.75	-58.53	-60.25	-19.41
Accidents	111	?	?	?	-33.19	-19.09
Retail	135	?	?	?	-10.94	-10.68
Pumsb Star	163	?	?	?	-30.79	-18.25
DNA	180	?	?	?	-87.68	-14.75
MSWeb	294	-9.77	-9.81	-10.21	-10.19	-9.25
Book	500	-35.62	-36.02	-34.22	-37.83	-17.38
EachMovie	500	?	?	?	-64.83	-14.89
WebKB	839	-161.30	-159.85	-156.84	-163.43	-15.89
Kosarak	883	?	?	?	-19.87	-9.03
Reuters-52	889	-89.54	-89.27	-91.23	-94.41	-17.58
20 NG	910	-159.56	-159.65	-156.77	-164.13	-18.91
BBC	1058	?	?	?	-261.80	-14.22
Ad	1556	?	?	?	-16.41	-10.31

does not appear in the dataset (training or testing) is an impossible example? Under this assumption, the domain constraints correspond to a sentence which is obtained by disjoining all examples in the dataset. This is indeed what we did in our last experiment. We compiled the mentioned disjunction (which is a DNF) into an SDD, learned its parameters using the training dataset, and measured the quality of learned PSDD using the testing dataset.

We compared the obtained results on the corpus in Table 4; see, e.g., (Davis and Domingos 2010) for more statistics on these datasets. These real-world datasets are commonly used to evaluate algorithms for learning probabilistic graphical models as we just discussed.⁵ Here, ACMN are Arithmetic Circuits for Bayesian networks, ACBN are Arithmetic Circuits for Markov networks, LTM are Latent Tree Models, and CLT are Chow-Liu trees. Each algorithm learns its corresponding model using the training dataset, and evaluates it using the testing dataset. The likelihoods of learned models are given in Table 4; these likelihoods can also be found in (Lowd and Rooshenas 2013). For the two datasets over a small number of variables, our PSDDs obtain competitive likelihoods (smaller magnitude is better). In datasets over a larger number of variables, the PSDDs obtain better, and often significantly better, likelihoods than other learned models.

While this experiment is contrary to what is usual, it is both revealing and suggestive. First, it highlights the potential impact of accommodating logical constraints into the learning process. Second, it begs the question: What if the data was indeed indicative of the domain constraints? For

⁵The models that were learned in Table 4 correspond to restricted Bayesian or Markov networks (they have a corresponding arithmetic circuit (AC) or have a tree structure). These restrictions are meant to allow one to compute likelihoods efficiently.

example, when the dataset is large enough, and the domain is known to be constrained enough, it is not unreasonable to assume that the dataset is indicative of domain constraints. Moreover, theoretically, this assumption does hold in the limit. As far as we know, there is no existing approach that can exploit this assumption as done by this work. In fact, this observation and the need to address it is what prompted us to develop PSDDs in the first place.

Related Work

The PSDD can be viewed from two angles: As a knowledge representation and reasoning formalism, and as a learning formalism. We next discuss related work across both angles.

As a KR formalism, the PSDD is related to work on extending Binary Decision Diagrams (BDDs) to represent real-valued functions. The relationship between SDDs (which underly PSDDs) and BDDs is known (Darwiche 2011; Xue, Choi, and Darwiche 2012; Choi and Darwiche 2013). In sum though, BDDs branch on literals, instead of sentences, which leads to limitations that are inherited by their extensions (Xue, Choi, and Darwiche 2012). Algebraic Decision Diagrams (ADDs) (R.I. Bahar et al. 1993) in particular are often used to represent joint probability distributions. However, they are not factorized representations and, as such, they need to explicitly represent every distinct probability in a distribution. Affine ADDs (Sanner and McAllester 2005) and Edge-valued Binary Decision Diagrams (Lai and Sastry 1992) do offer a factorized representation, but their parameters are not interpretable as in PSDDs (i.e., as probabilities of the global distribution).

There is also a body of work on facilitating the representation of logical constraints in the context of probabilistic representations. This includes Richardson and Domingos (2006), who facilitate the representation of logical constraints in Markov networks. Additionally, Mateescu and Dechter (2008) separate an unconstrained probabilistic model (e.g., Bayesian network) from the logical constraints and define the target distribution as the former conditioned on the latter. Hence, individual parameters in their framework have no local semantics in the target distribution. Several more first-order probabilistic languages can directly express logical dependencies (De Raedt et al. 2008; Getoor and Taskar 2007). Within these representations, there has been considerable interest in probabilistic reasoning in the presence of logical constraints. See for example Poon and Domingos (2006) and Gogate and Dechter (2007).

Perhaps the most strongly related formalism is the probabilistic decision graph (Bozga and Maler 1999; Jaeger 2004). These circuits are governed by a variable forest instead of a vtree. Their parameters represent contextual conditional probabilities. However, for a fixed variable forest, probabilistic decision graphs are not a complete representation. There are distributions that cannot be represented because the variable forest already encodes certain conditional independencies. Hence, these representations are canonical only in a weaker sense although they have been learned from data (Jaeger, Nielsen, and Silander 2006). As a logical representation, they are situated in between BDDs and SDDs:

They provide stronger decompositions than BDDs but still branch over single variables instead of sentences.

Viewing PSDDs as a learning formalism brings up two potential connections to existing work. First, as a formalism that aims at learning *generative* probabilistic models with well defined bases (i.e., logical constraints), the PSDD is somewhat unique as this learning problem has not been posed this explicitly before. The only possible exception is (Chang et al. 2008), which learns a *discriminative* probabilistic model given logical constraints (they propose a generalization of linear models called Constrained Conditional Models).⁶ Second, as a formalism that aims at learning *tractable* representations of distributions, the PSDD falls into the recent body of work on *deep learning*. The closest connection here is to the sum-product network (Poon and Domingos 2011), which is a new (but also tractable) class of deep architectures. These networks represent probability distributions as a (deep) hierarchy of mixtures (essentially, an arithmetic circuit with latent variables). However, they do not take logical constraints as an input. Moreover, even though they provide a less constrained representation of distributions, compared to the PSDD, the implication is that learning becomes harder (e.g., there is no closed form for parameter learning under complete data).

Conclusion

We presented the PSDD as a representation of probability distributions that respect a given propositional theory. The PSDD is a complete and canonical representation, with parameters that are interpretable as conditional probabilities. The PSDD encodes context-specific independences, which can be derived from its structure. The PSDD is a tractable representation, allowing one to compute the probability of any term in time linear in its size. The PSDD has unique maximum likelihood parameters under complete data, which can be learned efficiently using a closed form. Preliminary experimental results suggest that the PSDD can be quite effective in learning distributions under domain constraints.

Acknowledgments

This work has been partially supported by ONR grant #N00014-12-1-0423, NSF grant #IIS-1118122, and the Research Foundation-Flanders (FWO-Vlaanderen).

Appendix

Completeness and Canonicity of PSDDs

We will now address two connected questions with regards to the representational power of PSDDs. The first question asks whether every distribution can be represented by a PSDD. The second question asks whether the PSDD representation is canonical (given a vtree). The answers to both

⁶Generative learning is concerned with learning a distribution that is optimized for generating the data, yet is not connected to any particular reasoning task. Discriminative learning is concerned with learning a distribution that is optimized for a classification task (i.e., a probabilistic classifier).

questions rely on the notion of *probabilistic XY-partitions*, which are based on the *XY-partitions* underlying SDDs.

We review the latter notion first. Consider a propositional sentence α over disjoint variables \mathbf{X} and \mathbf{Y} . One can always express this sentence as

$$\alpha = (p_1 \wedge s_1) \vee \dots \vee (p_k \wedge s_k) \quad (4)$$

where p_1, \dots, p_k mention only variables \mathbf{X} and s_1, \dots, s_k mention only variables \mathbf{Y} . One can also always ensure that each p_i is consistent, that $p_i \wedge p_j$ is inconsistent for $i \neq j$, and that $p_1 \vee \dots \vee p_k$ is valid. Under these conditions, the form in (4) is called an *XY-partition* of α (Darwiche 2011). One obtains an SDD for some sentence α by recursively decomposing it into *XY-partitions* (Darwiche 2011). A parallel notion exists for distributions.

Definition 4 Let $\Pr(\mathbf{XY})$ be a distribution with base α . A *probabilistic XY-partition* of this distribution has the form

$$(\Pr_{p_1}(\mathbf{X}), \Pr_{s_1}(\mathbf{Y}), \theta_1), \dots, (\Pr_{p_k}(\mathbf{X}), \Pr_{s_k}(\mathbf{Y}), \theta_k)$$

such that

- $(p_1 \wedge s_1) \vee \dots \vee (p_k \wedge s_k)$ is an *XY-partition* of base α .
- Each \Pr_{p_i} is a (prime) distribution with base p_i .
- Each \Pr_{s_i} is a (sub) distribution with base s_i .
- $\theta_i \geq 0$, $\theta_i = 0$ iff $s_i = \perp$, and $\theta_1 + \dots + \theta_k = 1$.
- $\Pr(\mathbf{xy}) = \Pr_{p_i}(\mathbf{x}) \cdot \Pr_{s_i}(\mathbf{y}) \cdot \theta_i$ where p_i satisfies $\mathbf{x} \models p_i$.

The above definition shows how one can decompose a distribution $\Pr(\mathbf{XY})$ into a number of smaller distributions $\Pr(\mathbf{X})$ and $\Pr(\mathbf{Y})$. More importantly though, the bases of these smaller distributions correspond to a decomposition of the base for distribution $\Pr(\mathbf{XY})$.

One can always decompose a distribution this way.

Theorem 9 Let \Pr be a probability distribution and suppose that its variables are partitioned into \mathbf{X} and \mathbf{Y} . There must exist a *probabilistic XY-partition* of distribution \Pr .

Corollary 3 Every probability distribution can be represented by a PSDD.

Let v be a vtree over the distribution variables. The root of this vtree partitions the variables into \mathbf{X} (variables in left subtree of v) and \mathbf{Y} (variables in right subtree of v). By Theorem 9, we can construct a *probabilistic XY-partition* of the distribution, which defines the root PSDD node. By repeating the process recursively, we can construct PSDDs for the prime and sub distributions of each constructed *XY-partition*, until we reach distributions over single variables.

A distribution may have multiple (or even many) *probabilistic XY-partitions* for the same sets of variables \mathbf{X} and \mathbf{Y} . However, exactly one of these is *compressed*.

Theorem 10 A *probabilistic XY-partition* is *compressed* iff its sub distributions are distinct. A distribution has a *unique compressed XY-partition* for each \mathbf{X} and \mathbf{Y} .⁷

⁷Consider a prime distribution \Pr_{p_i} where sub $s_i = \perp$. Since $\theta_i = 0$ in this case, the non-zero probabilities of the prime distribution \Pr_{p_i} are irrelevant to the distribution represented by the *probabilistic XY-partition*. The uniqueness claimed in this theorem is modulo such prime distributions.

A decision PSDD node is compressed iff its subs (not their bases) are distinct. A PSDD is compressed iff all its decision nodes are compressed. We now have our canonicity result.

Corollary 4 *Every distribution is induced by a unique, compressed PSDD (given a vtree).*

We close this section by stressing the following point. The notion of compression in PSDDs generalizes a corresponding notion of compression for SDDs (Darwiche 2011). In particular, an SDD is said to be compressed if every decision node has distinct subs. The main observation here is that the SDD of a compressed PSDD may itself be uncompressed. That is, a compressed PSDD may have two distinct sub distributions with equal bases.⁸ In fact, the PSDD representation is complete only if one allows the underlying SDDs to be uncompressed.

References

- Boutilier, C.; Friedman, N.; Goldszmidt, M.; and Koller, D. 1996. Context-specific independence in Bayesian networks. In *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (UAI)*, 115–123.
- Bozga, M., and Maler, O. 1999. On the representation of probabilities over structured domains. In *Computer Aided Verification*, 261–273. Springer.
- Chang, M.-W.; Ratinov, L.-A.; Rizzolo, N.; and Roth, D. 2008. Learning and inference with constraints. In Fox, D., and Gomes, C. P., eds., *AAAI*, 1513–1518. AAAI Press.
- Chavira, M.; Darwiche, A.; and Jaeger, M. 2006. Compiling relational Bayesian networks for exact inference. *International Journal of Approximate Reasoning* 42(1-2):4–20.
- Choi, A., and Darwiche, A. 2013. Dynamic minimization of sentential decision diagrams. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 187–194.
- Darwiche, A. 2002. A logical approach to factoring belief networks. In *Proceedings of KR*, 409–420.
- Darwiche, A. 2009. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press.
- Darwiche, A. 2011. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI*, 819–826.
- Davis, J., and Domingos, P. 2010. Bottom-up learning of Markov network structure. In *ICML*, 271–278.
- De Raedt, L.; Frasconi, P.; Kersting, K.; and Muggleton, S., eds. 2008. *Probabilistic inductive logic programming: theory and applications*. Berlin, Heidelberg: Springer-Verlag.
- Fierens, D.; Van den Broeck, G.; Thon, I.; Gutmann, B.; and De Raedt, L. 2011. Inference in probabilistic logic programs using weighted CNF's. In *Proceedings of UAI*, 211–220.
- Getoor, L., and Taskar, B., eds. 2007. *An Introduction to Statistical Relational Learning*. MIT Press.
- Gogate, V., and Dechter, R. 2007. Samplesearch: A scheme that searches for consistent samples. In *International Conference on Artificial Intelligence and Statistics*, 147–154.
- Halpern, J. 1990. An analysis of first-order logics of probability. *Artificial Intelligence* 46(3):311–350.
- Jaeger, M.; Nielsen, J. D.; and Silander, T. 2006. Learning probabilistic decision graphs. *International Journal of Approximate Reasoning* 42(1):84–100.
- Jaeger, M. 2004. Probabilistic decision graphs combining verification and ai techniques for probabilistic inference. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 12(supp01):19–42.
- Lai, Y.-T., and Sastry, S. 1992. Edge-valued binary decision diagrams for multi-level hierarchical verification. In *Proceedings of the 29th ACM/IEEE Design Automation Conference*, 608–613. IEEE Computer Society Press.
- Lowd, D., and Rooshenas, A. 2013. Learning Markov networks with arithmetic circuits. In *AISTATS*, 406–414.
- Mateescu, R., and Dechter, R. 2008. Mixed deterministic and probabilistic networks. *Annals of Mathematics and Artificial Intelligence* 54(1-3):3–51.
- Nilsson, N. 1986. Probabilistic logic. *Artificial intelligence* 28(1):71–87.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, California.
- Pipatsrisawat, K., and Darwiche, A. 2008. New compilation languages based on structured decomposability. In *AAAI*, 517–522.
- Poole, D. 2003. First-order probabilistic inference. In *Proceedings of IJCAI*, 985–991.
- Poon, H., and Domingos, P. 2006. Sound and efficient inference with probabilistic and deterministic dependencies. In *AAAI*, volume 6, 458–463.
- Poon, H., and Domingos, P. 2011. Sum-product networks: A new deep architecture. *Proc. 12th Conf. on Uncertainty in Artificial Intelligence* 337–346.
- R.I. Bahar; E.A. Frohm; C.M. Gaona; G.D. Hachtel; E. Macii; A. Pardo; and F. Somenzi. 1993. Algebraic Decision Diagrams and Their Applications. In *IEEE /ACM International Conference on CAD*, 188–191. Santa Clara, California: IEEE Computer Society Press.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1):107–136.
- Sanner, S., and McAllester, D. A. 2005. Affine algebraic decision diagrams (aadds) and their application to structured probabilistic inference. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 1384–1390.
- Van den Broeck, G.; Taghipour, N.; Meert, W.; Davis, J.; and De Raedt, L. 2011. Lifted probabilistic inference by first-order knowledge compilation. In *Proceedings of IJCAI*, 2178–2185.
- Xue, Y.; Choi, A.; and Darwiche, A. 2012. Basing decisions on sentences in decision diagrams. In *AAAI*, 842–849.

⁸This is never true in our experiments since we start off with a compressed SDD (the SDD package outputs compressed SDDs).