

First-Order Context-Specific Likelihood Weighting in Hybrid Probabilistic Logic Programs

Nitesh Kumar

*Department of Computer Science
KU Leuven, Belgium*

NITESH.KR369@GMAIL.COM

Ondřej Kuželka

*Faculty of Electrical Engineering
Czech Technical University in Prague, Czech Republic*

ONDREJ.KUZELKA@FEL.CVUT.CZ

Luc De Raedt

*Department of Computer Science
KU Leuven, Belgium*

LUC.DERAEDT@KULEUVEN.BE

Abstract

Statistical relational AI and probabilistic logic programming have so far mostly focused on discrete probabilistic models. The reasons for this is that one needs to provide constructs to succinctly model the independencies in such models, and also provide efficient inference.

Three types of independencies are important to represent and exploit for scalable inference in hybrid models: conditional independencies elegantly modeled in Bayesian networks, context-specific independencies naturally represented by logical rules, and independencies amongst attributes of related objects in relational models succinctly expressed by combining rules.

This paper introduces a hybrid probabilistic logic programming language, DC#, which integrates distributional clauses' syntax and semantics principles of Bayesian logic programs. It represents the three types of independencies qualitatively. More importantly, we also introduce the scalable inference algorithm FO-CS-LW for DC#. FO-CS-LW is a first-order extension of the context-specific likelihood weighting algorithm (CS-LW), a novel sampling method that exploits conditional independencies and context-specific independencies in ground models. The FO-CS-LW algorithm upgrades CS-LW with unification and combining rules to the first-order case.

1. Introduction

Statistical relational AI (StarAI) and probabilistic logic programming (PLP) (De Raedt & Kimmig, 2015; De Raedt, et al., 2016) have contributed many languages for declaratively modeling expressive probabilistic models and have devised numerous inference techniques. They have been applied to many applications in databases, knowledge graphs, social networks, robotics, chemical compounds, genomics, etc.

To enable scalable probabilistic inference, it is essential to represent three different types of independencies in the modeling language. Firstly, the traditional classical conditional independencies (CIs) represented by Bayesian networks (BNs). Secondly, the context-specific independencies (CSIs): independencies that hold only in certain contexts (Boutilier, et al., 1996). These independencies arise due to structures present in conditional probability distributions (CPDs) of BNs, which BNs do not qualitatively represent, but rule-based rep-

representations do by making structures explicit in the clauses (Poole, 1997). Thirdly, the combining rules such as NoisyOR to express probabilistic influences among attributes of related objects (Koller, et al., 2007; Jaeger, 2007). Combining rules are particularly interesting since they allow one to qualitatively represent independence of causal influences (Zhang & Poole, 1996, ICIs), where each influence is considered independent of others. This independence is natural and commonly assumed to keep relational models succinct (Natarajan, et al., 2008). In probabilistic logic programs that deal with only discrete random variables, combining rules are the core component (Kersting & De Raedt, 2007; De Raedt, Kimmig, & Toivonen, 2007).

Over the past few decades, many PLP languages have been proposed; however, only a few of them are hybrid, i.e., support both discrete and continuous random variables. Nevertheless, such hybrid PLP are needed to cope with applications in areas such as activity recognition, robotics, sensing, perception, etc. Current hybrid PLP languages suffer from two problems. Firstly, they generally do not support combining rules (Gutmann, Jaeger, & De Raedt, 2010; Gutmann, et al., 2011; Islam, Ramakrishnan, & Ramakrishnan, 2012; Michels, Hommersom, & Lucas, 2016; Alberti, et al., 2017). Secondly, and more importantly, there exist, to the best of our knowledge, no inference techniques that exploit all three types of independencies for hybrid PLPs.

To remedy this, we first introduce DC# (pronounced “DC sharp”), a hybrid PLP language that supports combining rules. DC# uses a special form of clauses called distributional clauses (DCs) to express probabilistic knowledge. We borrow the syntax of DCs from (Nitti, De Laet, & De Raedt, 2016) but introduce an extended new semantics based on Bayesian Logic Programs (Kersting & De Raedt, 2007, BLPs). Thus, in terms of representation, DC#, a rule-based representation, differs from graphical model-based relational representations, such as BLPs, which are associated with CPDs. However, the semantics of DC# are based on BLPs, so DC# programs can be seen as BLPs qualitatively representing CSIs.

Our second contribution is the first-order context-specific likelihood weighting (FO-CS-LW) algorithm that exploits these three types of independencies for scalable inference in DC# programs. Before going to the first-order case, we introduce the CS-LW algorithm for ground programs. CS-LW exploits both CIs and CSIs, which is an approximate inference algorithm that, until our earlier work, has not been well-studied yet (Kumar & Kuželka, 2021).

There exist state-of-the-art inference algorithms for exact inference (Friedman & Van Den Broeck, 2018) in discrete models. These algorithms are based on the knowledge compilation technique (Darwiche & Marquis, 2002) that uses logical reasoning to exploit CSIs. However, for many models, exact inference quickly becomes infeasible. Stochastic sampling for approximate inference is a standard solution. Sampling algorithms are simple yet powerful tools, and they can be applied to arbitrary complex hybrid models, unlike exact inference. It is widely believed that CSI properties in distributions are difficult to exploit for approximate inference (Friedman & Van Den Broeck, 2018). To solve this difficult problem, we introduce the context-specific likelihood weighting (CS-LW), a sampling algorithm that exploits both CI and CSI properties, leading to faster convergence and faster sampling speed than standard likelihood weighting (LW).

Next, we extend CS-LW to first-order DC# programs specifying relational probabilistic models. Due to the use of combining rules, such models, when grounded, have many symmetrically repeated parameters. Inference algorithms designed for grounded models can not exploit these symmetries, rendering inference infeasible even for very simple relational models. It is widely believed that the inference can be feasible if one does not ground out these models, but reason at the first-order level with unification (Poole, 2003). There have been several attempts at doing this for various simple languages (Kisynski & Poole, 2009; Choi, Braz, & Bui, 2011; Van Den Broeck, et al., 2011; Beame, et al., 2015), but not for hybrid and expressive languages like DC#. For example, well-known graphical model-based relational representation languages that do not qualitatively represent CSIs, construct ground BNs for inference (Getoor, et al., 2007; Kersting & De Raedt, 2007). Similarly, well-known PLP systems, which do represent CSIs qualitatively, first ground the first-order programs and then perform inference at the ground level (Fierens, et al., 2015). In contrast, FO-CS-LW reasons at the first-order level. Using the tools of logic, i.e., unification, substitution, and resolution, FO-CS-LW samples only relevant variables from first-order DC# programs determined by various forms of independencies present in the programs. We empirically demonstrate that FO-CS-LW scales with domain size and provide an open-source implementation of our framework¹.

This paper is a significantly extended and completed version of our previous paper (Kumar & Kuželka, 2021), where we introduced CS-LW to exploit the structures present in CPDs of BNs. The present paper first introduces a language to describe first-order probabilistic models and then extends CS-LW to the first-order case.

Contribution We summarise our contributions in this paper as follows:

- We introduce a new PLP language DC# that supports combining rules to describe hybrid relational probabilistic models succinctly.
- We present a novel sampling methodology, CS-LW, that exploits both CIs and CSIs for probabilistic inference in BNs and ground probabilistic logic programs.
- We present a first-order extension of CS-LW that applies directly to first-order DC# programs and in addition exploits the symmetries present through combining rules.
- We empirically show that our inference algorithm scales with the domain size when applied to hybrid relational probabilistic models described as DC# programs.

Organization The paper is organized as follows. We start by motivating our discussion with some examples in Section 2. In Section 3, we review the standard likelihood weighting and basic concepts of logic programming. Section 4 presents the DC# language. Section 5 presents the CS-LW algorithm for ground DC# programs describing BNs, which is then extended to first-order DC# programs in Section 6. We then evaluate our algorithms in Section 7. Before concluding, we finally touch upon related work and directions for future work.

1. The code is publicly available: <https://github.com/niteshroyal/DC-Sharp>

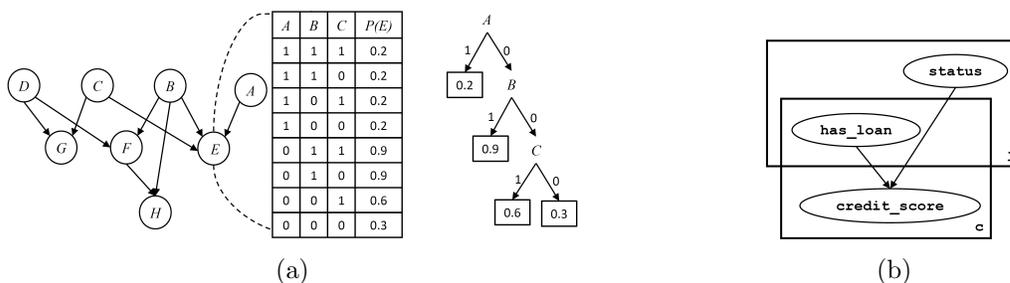


Figure 1: (a) Context-Specific Independence; (b) A Bayesian network with client and loan plates

2. Motivating Examples

Let us illustrate, with examples, the independencies that DC# programs will qualitatively represent and that our algorithm will exploit. Consider a BN in Figure 1a, where a tree-structure is present in the CPD of variable E . If one observes the CPD carefully, one can conclude that $P(E \mid A = 1, B, C) = P(E \mid A = 1)$, that is, $P(E \mid A = 1, B, C)$ is the same for all values of B and C . The variable E is said to be independent of variables $\{B, C\}$ in the *context* $A = 1$. This local independence statement corresponds to the influence of edges $\{B \rightarrow E, C \rightarrow E\}$ vanishing in this context; consequently, it may have global implications. For example, $E \perp B, C \mid A = 1$ implies $E \perp D \mid H, A = 1$. These independencies are called CSIs. They arise naturally in various real-world situations (Poole & Zhang, 2003), including when one writes *if-then conditions* in probabilistic programs (Li & Russell, 2013). Our CS-LW algorithm aims to exploit CSIs arising due to the structures present within CPDs of BNs when ground DC# programs describe such BNs.

The exploitation of CSIs in relational probabilistic models is even more crucial since a huge amount of CSIs is present in these models. As a simple example, consider the model in Figure 1b, where the plate notation is used to represent direct influence relationships among credit scores of c clients, statuses of 1 loans, and $c \times 1$ number of `has_loan` type random variables for each client-loan pair, which can be either true or false. Given that a client has a loan, it is easy to imagine that the status of the loan may affect the client’s credit score. On the other hand, if a client has only a few loans, then it is just as easy to imagine that the status of the loans that the client does not have will not affect the client’s credit score. That is, the client’s credit is independent of the status of all those loans that the client does not have. This is a first-order level CSI, which BNs with plates do not qualitatively represent. We will introduce a PLP language to represent it.

Furthermore, it is natural to imagine that the client’s credit score depends on the number of loans that the client has with approved and rejected statuses but not on the identity of those approved and rejected loans. That is, loans are *exchangeable* objects. In such a case, it is common to consider that each loan’s status independently has its own probabilistic influence on the client’s credit score, and the final probabilistic influence is a combination of all these influences (Natarajan et al., 2008; Koller & Friedman, 2009). This independence assumption implies exchangeability. Relational models written as DC# programs quali-

tatively represent these independencies and the first-order level CSIs that our FO-CS-LW algorithm aims to exploit for scalability.

3. Background

A Bayesian network \mathcal{B} is a pair $(\mathcal{G}, \mathcal{D})$, where \mathcal{G} is a directed acyclic graph structure specifying *direct influence* relationships among random variables (nodes), and \mathcal{D} is a set of CPDs associated with each node. The CPD specifies the conditional probability of the variable given its parents. The graph structure \mathcal{G} represents local CI statements, which states that each variable is conditionally independent of its non-descendants given its parents. The local CIs and the set of CPDs \mathcal{D} together induce a joint probability distribution P over all the variables.

We denote random variables (RVs) with uppercase letters (A) and their assignments with lowercase letters (a). Bold letters denote sets of RVs (\mathbf{A}) and their assignments (\mathbf{a}). Parents of the variable A are denoted with $\mathbf{Pa}(A)$ and their assignments with $\mathbf{pa}(A)$. Suppose P is a probability distribution over disjoint sets of variables $\mathbf{E}, \mathbf{X}, \mathbf{Z}$, then \mathbf{E} denotes a set of observed variables, \mathbf{X} a set of unobserved query variables and \mathbf{Z} a set of unobserved variable other than query variables. The expected value of A relative to a distribution Q is denoted by $\mathbb{E}_Q[A]$.

3.1 Likelihood Weighting

Next, we briefly review likelihood weighting (LW), one of the most popular sampling algorithms for BNs.

A typical query to the distribution $P(\mathbf{E}, \mathbf{X}, \mathbf{Z})$ is to compute $P(\mathbf{x}_q \mid \mathbf{e})$, that is, the probability of \mathbf{X} being assigned \mathbf{x}_q given that \mathbf{E} is assigned \mathbf{e} . Following Bayes's rule, we have:

$$P(\mathbf{x}_q \mid \mathbf{e}) = \frac{P(\mathbf{x}_q, \mathbf{e})}{P(\mathbf{e})} = \frac{\sum_{\mathbf{x}, \mathbf{z}} P(\mathbf{x}, \mathbf{z}, \mathbf{e}) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}} P(\mathbf{x}, \mathbf{z}, \mathbf{e})} = \mu,$$

where $f(\mathbf{x})$ is an indicator function $\mathbb{1}\{\mathbf{x} = \mathbf{x}_q\}$, which takes value 1 when $\mathbf{x} = \mathbf{x}_q$, and 0 otherwise. We can estimate μ using LW if we specify P using a Bayesian network. LW belongs to a family of importance sampling schemes that are based on the observation,

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}} Q(\mathbf{x}, \mathbf{z}, \mathbf{e}) f(\mathbf{x}) (P(\mathbf{x}, \mathbf{z}, \mathbf{e}) / Q(\mathbf{x}, \mathbf{z}, \mathbf{e}))}{\sum_{\mathbf{x}, \mathbf{z}} Q(\mathbf{x}, \mathbf{z}, \mathbf{e}) (P(\mathbf{x}, \mathbf{z}, \mathbf{e}) / Q(\mathbf{x}, \mathbf{z}, \mathbf{e}))}, \tag{1}$$

where Q is a *proposal distribution* such that $Q > 0$ whenever $P > 0$. The distribution Q is different from P and is used to draw independent samples. Generally, Q is selected such that the samples can be drawn easily. In the case of LW, to draw a sample, variables $X_i \in \mathbf{X} \cup \mathbf{Z}$ are assigned values drawn from $P(X_i \mid \mathbf{pa}(X_i))$ and variables in \mathbf{E} are assigned their observed values. These variables are assigned in a topological ordering relative to the graph structure of \mathcal{B} . Thus, the proposal distribution in the case of LW can be described as follows:

$$Q(\mathbf{X}, \mathbf{Z}, \mathbf{E}) = \prod_{X_i \in \mathbf{X} \cup \mathbf{Z}} P(X_i \mid \mathbf{Pa}(X_i)) \mid_{\mathbf{E}=\mathbf{e}}.$$

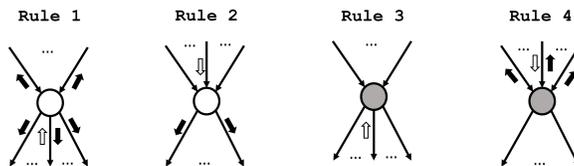


Figure 2: The four rules of Bayes-ball algorithm that decide next visits (indicated using \blackuparrow) based on the direction of the current visit (indicated using \uparrow) and the type of variable. To distinguish observed variables from unobserved variables, the former type of variables are shaded.

Consequently, it is easy to compute the *likelihood ratio* $P(\mathbf{x}, \mathbf{z}, \mathbf{e})/Q(\mathbf{x}, \mathbf{z}, \mathbf{e})$ in Equation 1. All factors in the numerator and denominator of the fraction cancel out except for $P(x_i | \mathbf{pa}(X_i))$ where $x_i \in \mathbf{e}$. Thus,

$$\frac{P(\mathbf{X}, \mathbf{Z}, \mathbf{e})}{Q(\mathbf{X}, \mathbf{Z}, \mathbf{e})} = \prod_{x_i \in \mathbf{e}} P(x_i | \mathbf{Pa}(X_i)) = \prod_{x_i \in \mathbf{e}} W_{x_i} = W_{\mathbf{e}},$$

where W_{x_i} , which is also a RV, is the *weight of evidence* x_i . The *likelihood ratio* $W_{\mathbf{e}}$ is the product of all of these weights, and thus, it is also a RV. Given M independent weighted samples from $Q(\mathbf{X}, \mathbf{Z}, \mathbf{E})$, we can estimate the query:

$$\hat{\mu} = \frac{\sum_{m=1}^M f(\mathbf{x}[m])w_{\mathbf{e}}[m]}{\sum_{m=1}^M w_{\mathbf{e}}[m]}. \tag{2}$$

3.2 Likelihood Weighting + Bayes-Ball

In the previous section, we used all random variables to estimate μ . However, due to CIs encoded by the graph structure in a Bayesian network \mathcal{B} , observed states and CPDs of only some variables “might” be required for computing μ . These variables are called *requisite variables*. To get a better estimate of μ , it is recommended to use only these variables. The standard approach is to first apply the Bayes-ball algorithm (Shachter, 1998) over the graph to obtain a sub-network of requisite variables, then simulate the sub-network to obtain the weighted samples. An alternative approach is to use Bayes-ball to simulate the original network \mathcal{B} and focus on only requisite variables to obtain the weighted samples. This approach is trivial and might already be used by many BN inference tools. However, it is not described in the literature clearly, so we will describe it next. It will form the basis of our discussion on CS-LW, where we will also exploit structures within CPDs of BNs.

To obtain the samples, we need to traverse the graph in a topological ordering. The Bayes-ball algorithm, which is linear in the graph’s size, can be used for it. The advantage of using Bayes-ball is that it also detects CIs; thus, it traverses only a sub-graph that depends on the query and evidence. We can also keep assigning unobserved variables, and weighting observed variables along with traversing the graph. In this way, we assign/weight only requisite variables. The Bayes-ball algorithm uses four rules to traverse the graph (when deterministic variables are absent in \mathcal{B}), and marks variables to avoid repeating the

same action. These rules are illustrated in Figure 2. Next, we discuss these rules and also indicate how to assign/weigh variables, resulting in a new algorithm that we call *Bayes-ball simulation of BNs*. Starting with all query variables scheduled to be visited as if from one of their children, we apply the following rules until no more variables can be visited:

1. When the visit of an unobserved variable $U \in \mathbf{X} \cup \mathbf{Z}$ is from a child, and U is not marked on top, then do these in the order: i) Mark U on top; ii) Visit all its parents; iii) Sample a value y from $P(U \mid \mathbf{pa}(U))$ and assign y to U ; iv) If U is not marked on bottom, then mark U on bottom and visit all its children.
2. When the visit of an unobserved variable is from a parent, and the variable is not marked on bottom, then mark the variable on bottom and visit all its children.
3. When the visit of an observed variable is from a child, then do nothing.
4. When the visit of an observed variable $E \in \mathbf{E}$ is from a parent, and E is not marked on top, then do these in the order: i) Mark E on top; ii) Visit all its parents; iii) Let e be an observed value of E and let w be the probability at e according to $P(E \mid \mathbf{pa}(E))$, then the weight of E is w .

The above rules define an order for visiting parents and children so that variables are assigned/weighted in a topological ordering. Indeed we can define the order since the original rules for Bayes-ball do not prescribe any order. The marks record important information, and the following result holds.

Lemma 1. *Let $\mathbf{E}_\star \subseteq \mathbf{E}$ be marked on top, $\mathbf{E}_\star \subseteq \mathbf{E}$ be visited but not marked on top, and $\mathbf{Z}_\star \subseteq \mathbf{Z}$ be marked on top. Then the query μ can be computed as follows,*

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star)} \quad (3)$$

The proof is straightforward and is present in Appendix A. Now, since $\mathbf{X}, \mathbf{Z}_\star, \mathbf{E}_\star, \mathbf{E}_\star$ are variables of \mathcal{B} and they form a sub-network \mathcal{B}_\star such that variables in \mathbf{E}_\star do not have any parent, we can write,

$$P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star) = \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\star \cup \mathbf{e}_\star} P(u_i \mid \mathbf{pa}(U_i))$$

such that $\forall p \in \mathbf{pa}(U_i) : p \in \mathbf{x} \cup \mathbf{z}_\star \cup \mathbf{e}_\star \cup \mathbf{e}_\star$. This means the CPDs of some observed variables are not required for computing μ . Now we define these variables.

Definition 1. *The observed variables whose observed states and CPDs might be required to compute μ will be called diagnostic evidence.*

Definition 2. *The observed variables whose observed states, but not their CPDs, might be required to compute μ will be called predictive evidence.*

Diagnostic evidence (denoted by \mathbf{e}_\star) is marked on top, while predictive evidence (denoted by \mathbf{e}_\star) is visited but not marked on top. The variables $\mathbf{X}, \mathbf{Z}_\star, \mathbf{E}_\star, \mathbf{E}_\star$ will be called requisite variables.

Example 1. Consider the network of Figure 1a, and assume that our evidence is $\{D = 1, F = 1, G = 0, H = 1\}$, and our query is $\{E = 0\}$. Suppose we start by visiting the query variable from its child and apply the four rules of Bayes-ball. One can easily verify that observed variables F, G, H will be marked on top; hence $\{F = 1, G = 0, H = 1\}$ is diagnostic evidence (\mathbf{e}_\star). The observed variable D will only be visited; hence $\{D = 1\}$ is predictive evidence (\mathbf{e}_\star). Variables A, B, C, E will be marked on top and are requisite unobserved variables ($\mathbf{X} \cup \mathbf{Z}_\star$).

Now, we can sample from a factor Q_\star of Q such that,

$$Q_\star(\mathbf{X}, \mathbf{Z}_\star, \mathbf{E}_\star \mid \mathbf{E}_\star) = \prod_{X_i \in \mathbf{X} \cup \mathbf{Z}_\star} P(X_i \mid \mathbf{Pa}(X_i)) \mid_{\mathbf{E}_\star = \mathbf{e}_\star} \quad (4)$$

When we use Bayes-ball, precisely this factor is considered for sampling. Starting by first setting \mathbf{E}_\star to their observed values, $\mathbf{X} \cup \mathbf{Z}_\star$ is assigned and \mathbf{e}_\star is weighted in the topological ordering. Given M weighted samples $\mathcal{D}_\star = \langle \mathbf{x}[1], w_{\mathbf{e}_\star}[1] \rangle, \dots, \langle \mathbf{x}[M], w_{\mathbf{e}_\star}[M] \rangle$ from Q_\star , we can estimate:

$$\tilde{\mu} = \frac{\sum_{m=1}^M f(\mathbf{x}[m]) w_{\mathbf{e}_\star}[m]}{\sum_{m=1}^M w_{\mathbf{e}_\star}[m]}. \quad (5)$$

In this way, we sample from a lower-dimensional space; thus, the new estimator $\tilde{\mu}$ has a lower variance compared to $\hat{\mu}$ due to the Rao-Blackwell theorem. Consequently, fewer samples are needed to achieve the same accuracy. Hence, for improved inference, we exploit CIs encoded by the graph structure in \mathcal{B} .

3.3 Context-Specific Independence

Next, we formally define the independencies that arise due to the structures present within CPDs, which were informally discussed in Section 2.

Definition 3. Let P be a probability distribution over variables \mathbf{U} , and let $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ be disjoint subsets of \mathbf{U} . The variables \mathbf{A} and \mathbf{B} are independent given \mathbf{D} and context \mathbf{c} if $P(\mathbf{A} \mid \mathbf{B}, \mathbf{D}, \mathbf{c}) = P(\mathbf{A} \mid \mathbf{D}, \mathbf{c})$ whenever $P(\mathbf{B}, \mathbf{D}, \mathbf{c}) > 0$. This is denoted by $\mathbf{A} \perp \mathbf{B} \mid \mathbf{D}, \mathbf{c}$. If \mathbf{D} is empty then \mathbf{A} and \mathbf{B} are independent given context \mathbf{c} , denoted by $\mathbf{A} \perp \mathbf{B} \mid \mathbf{c}$.

Independence statements of the above form are called *context-specific independencies* (CSIs). When \mathbf{A} is independent of \mathbf{B} given all possible assignments to \mathbf{C} then we have: $\mathbf{A} \perp \mathbf{B} \mid \mathbf{C}$. The independence statements of this form are generally referred to as *conditional independencies* (CIs). Thus, CSI is a more fine-grained notion than CI. The graphical structure in \mathcal{B} can only represent CIs. Any CI can be verified in linear time in the size of the graph. However, verifying any arbitrary CSI has been recently shown to be coNP-hard (Corander, et al., 2019).

3.4 A Bit of Logic Programming

Probabilistic logic programming is a probabilistic characterization of logic programming. So, before describing our system, in this section, we review relevant syntactic and semantic notions related to logic programming. More details can be found in (Nilsson & Małuszynski, 1995).

An *atom* $p(\mathbf{t}_1, \dots, \mathbf{t}_n)$ consists of a predicate p/n of arity n and terms $\mathbf{t}_1, \dots, \mathbf{t}_n$. A *term* is either a constant (written in lowercase), a variable (in uppercase), or a structured term of the form $f(\mathbf{u}_1, \dots, \mathbf{u}_k)$ where f is a functor and the \mathbf{u}_i are terms. For example, `has_account(ann, L)`, `has_account(ann, a.1)` and `has_account(ann, func(A))` are atoms and `ann`, `L`, `a.1` and `func(A)` are terms. A *literal* is an atom or the negation of an atom. A *positive literal* is an atom. A *negative literal* is the negation of an atom. A *clause* is a universally quantified disjunction of literals. A *definite clause* is a clause which contains exactly one positive literal and zero or more negative literals. For example, $\forall(A_0 \vee \neg A_1 \vee \dots \vee \neg A_n)$ is a definite clause, where A_0, A_1, \dots, A_n are atoms. In logic programming, one usually writes definite clauses in the implication form $A_0 \leftarrow A_1, \dots, A_n$ (where we omit the universal quantifiers for ease of writing). Here, the atom A_0 is called *head* of the clause; and the set of atoms $\{A_1, \dots, A_n\}$ is called *body* of the clause. A clause with an empty body is called a *fact*. A *definite program* consists of a finite set of definite clauses.

Example 2. A clause $C \equiv \text{has_loan}(C, L) \leftarrow \text{has_account}(C, A), \text{account_loan}(A, L)$ is a definite clause. Intuitively, it states that a client C has a loan L if C has an account A and A is associated to the loan L .

An *expression*, which can be either a term, an atom or a clause, is *ground* if it does not contain any variable. A *substitution* $\theta = \{V_1/\mathbf{t}_1, \dots, V_m/\mathbf{t}_m\}$ assigns terms \mathbf{t}_i to variables V_i . The element V_i/\mathbf{t}_i is a *binding* for variable V_i . Applying θ to an expression \mathcal{E} yields $\mathcal{E}\theta$, the *instance* of \mathcal{E} , where all occurrences of V_i in \mathcal{E} are replaced by the corresponding terms \mathbf{t}_i . A substitution θ is a *grounding* for \mathcal{E} if $\mathcal{E}\theta$ is ground, i.e., contains no variables (when there is no risk of confusion we drop “for \mathcal{E} ”).

Example 3. Applying a substitution $\theta = \{C/\text{ann}\}$ to the clause C from Example 2 yields $C\theta$ which is `has_loan(ann, L) ← has_account(ann, A), account_loan(A, L)`.

A substitution θ *unifies* two expressions \mathcal{E}_1 and \mathcal{E}_2 if $\mathcal{E}_1\theta$ and $\mathcal{E}_2\theta$ are identical (denoted $\mathcal{E}_1\theta = \mathcal{E}_2\theta$). Such a substitution is called a *unifier*. Unifiers may not always exist. If there exists a unifier for two expressions \mathcal{E}_1 and \mathcal{E}_2 , we call such atoms *unifiable* and we say that \mathcal{E}_1 and \mathcal{E}_2 *unify*.

Example 4. A substitution $\theta = \{C/\text{ann}, M/l.1, L/l.1\}$ unify `has_loan(ann, L)` and `has_loan(C, M)`.

A substitution θ is said to be more general than a substitution σ iff there exists a substitution σ' such that $\sigma = \theta\sigma'$. A unifier θ is said to be a *most general unifier (mgu)* of two expressions iff θ is more general than any other unifier of the expressions. Expression \mathcal{E}_1 is a *renaming* of \mathcal{E}_2 if they differ only in the names of variables.

Example 5. Unifiers $\{C/\text{ann}, M/L\}$ and $\{C/\text{ann}, L/M\}$ are both most general unifiers of `has_loan(ann, L)` and `has_loan(C, M)`. The resulting applications,

$$\begin{aligned} \text{has_loan}(C, M)\{C/\text{ann}, M/L\} &= \text{has_loan}(\text{ann}, L) \\ \text{has_loan}(\text{ann}, L)\{C/\text{ann}, L/M\} &= \text{has_loan}(\text{ann}, M) \end{aligned}$$

are renamings of each other.

The *Herbrand universe* of a definite program \mathcal{P} , denoted $U_{\mathcal{P}}$, is the set of all ground terms constructed from functors and constants appearing in \mathcal{P} . The *Herbrand base* $B_{\mathcal{P}}$ is the set of all ground atoms that can be constructed by using predicates from \mathcal{P} with ground terms from $U_{\mathcal{P}}$ as arguments. Subsets of the Herbrand base are called *Herbrand interpretations*. A Herbrand interpretation \mathcal{I} is a model of a clause $A_0 \leftarrow A_1, \dots, A_n$ iff for all grounding substitutions θ , such that $\{A_1\theta, \dots, A_n\theta\} \subseteq \mathcal{I}$, it also holds that $A_0\theta \in \mathcal{I}$. A *Herbrand model* of a set of clauses is a Herbrand interpretation which is a model of every clause in the set.

The *least Herbrand model* of a definite program \mathcal{P} , denoted $M_{\mathcal{P}}$, is the intersection of all Herbrand models of \mathcal{P} , i.e., $M_{\mathcal{P}}$ is the set of all ground atoms that are logical consequences of the program. $M_{\mathcal{P}}$ is unique for definite programs and can be constructed by repeatedly applying the so-called $T_{\mathcal{P}}$ operator, which is defined as a function on Herbrand interpretations of \mathcal{P} as follows:

$$T_{\mathcal{P}}(\mathcal{I}) := \{A_0\theta \mid A_0 \leftarrow A_1, \dots, A_n \in \mathcal{P} \wedge \{A_1\theta, \dots, A_n\theta\} \subseteq \mathcal{I}\},$$

where, θ are grounding substitutions. Let \mathcal{I}_1 be the set of all ground facts in the program. Now, applying the operator on \mathcal{I}_1 , it is possible to use every ground instance of each clause to construct new ground atoms from \mathcal{I}_1 . In this way, a new set $\mathcal{I}_2 := T_{\mathcal{P}}(\mathcal{I}_1)$ is obtained, which can be used again to construct more ground atoms. The new atoms added to \mathcal{I}_{i+1} are those which must follow immediately from \mathcal{I}_i . It is possible to construct $M_{\mathcal{P}}$ by recursively applying the operator until a *fixpoint* is reached ($T_{\mathcal{P}}(M_{\mathcal{P}}) = M_{\mathcal{P}}$), i.e., until no more ground atoms can be constructed.

A *query* Q is of the form B_1, \dots, B_m where the B_j are atoms and all variables are understood to be existentially quantified. Given a definite program \mathcal{P} , a correct answer to the query Q is a substitution θ such that $Q\theta$ is entailed by \mathcal{P} , denoted by $\mathcal{P} \models Q\theta$. That is, $Q\theta$ belongs to $M_{\mathcal{P}}$. The answer substitution θ is often computed using *SLD-resolution*. Finally, the answer set of Q is the set of all correct answer substitutions to Q .

4. DC#: A Representation Language for Hybrid Relational Models

This section presents a PLP language DC# for describing hybrid relational probabilistic models. The syntax of our language is based on the elegant syntax of distributional clauses (DCs) used by (Nitti et al., 2016). However, we extend its semantics to support combining rules. The new semantics, however, do not allow for describing open-universe probabilistic models (OUPMs) (Milch & Russell, 2009), which was possible in the previous system. The semantics supporting both OUPMs and combining rules is another topic of research. We do not study it in this paper. The new semantics have been developed along the lines of Bayesian Logic Programs (BLPs) (Kersting & De Raedt, 2007).

4.1 Syntax

DC is a natural extension of definite clauses for representing conditional probability distributions.

Definition 4. A DC is a formula of the form $A_0 \sim D \leftarrow A_1, \dots, A_n$, where \sim is a special binary predicate used in infix notation, and $A_i (i > 0)$ are atoms. Term A_0 is called a random variable term and D is called a distributional term.

Intuitively, the clause *defines* that RV $A_0\theta$ is distributed as $D\theta$ whenever all $A_i\theta$ are true for a grounding substitution θ . The ground terms $A_0\theta$ and $D\theta$ belong to the Herbrand universe.

Ground RV terms are interpreted as RVs. To refer to the values of RV terms, we use a binary predicate \cong , which is used in infix notation for convenience. A ground atom $A\theta \cong x$ is defined to be true if x is the value of RV (or ground RV term) $A\theta$.

Example 6. *Consider the following clause,*

$$\text{credit_score}(C) \sim \text{gaussian}(755.5, 0.1) \leftarrow \text{has_loan}(C, L) \cong \text{true}, \text{status}(L) \cong \text{appr}.$$

*Applying a grounding substitution $\{C/\text{ann}, L/1.1\}$ to the clause results in defining a RV $\text{credit_score}(\text{ann})$ distributed as $\text{gaussian}(755.5, 0.1)$ when RVs $\text{has_loan}(\text{ann}, 1.1)$ and $\text{status}(1.1)$ take values **true** and **appr** (“approved”) respectively; that is, when atoms $\text{has_loan}(\text{ann}, 1.1) \cong \text{true}$ and $\text{status}(1.1) \cong \text{appr}$ are true.*

A DC without body is called a *probabilistic fact*, for example:

$$\text{age}(\text{bob}) \sim \text{gaussian}(40, 0.2).$$

This clause states that the age of **bob** is normally distributed with mean 40 and variance 0.2. It is also possible to define deterministic variables that take only one value with probability 1, e.g., to express our absolute certainty that the age of **bob** takes value 40, we can write,

$$\text{age}(\text{bob}) \sim \text{val}(40).$$

Hence, it is also possible to write the definite clause \mathcal{C} of Example 2 as a DC like this:

$$\text{has_loan}(C, L) \sim \text{val}(\text{true}) \leftarrow \text{has_account}(C, A) \cong \text{true}, \text{account_loan}(A, L) \cong \text{true}$$

DC also supports comparing the values of RVs with constants or with values of other RVs in the ground instance of clauses. This can be done using the binary infix predicates $=$, $<$, $>$, \geq , and \leq , which are especially useful while conditioning continuous RVs as illustrated by the following clause:

$$\text{credit_score}(C) \sim \text{gaussian}(645.5, 0.1) \leftarrow \text{age}(C) \cong X, X < 40.$$

This clause specifies the distribution of a client’s credit score when the client’s age is less than 40.

A distributional program \mathbb{P} consists of a set of distributional clauses.

Example 7. *The following program describes probabilistic influences among attributes of clients and related loans. Here, we have used **t** as a shorthand for **true**, **f** for **false**, **a** for **approved**, and **d** for **declined**.*

```

client(ann) ~ val(t).
loan(1.1) ~ val(t).
loan(1.2) ~ val(t).
has_loan(C, L) ~ bernoulli(0.2) ← client(C) ≅ t, loan(L) ≅ t.
status(L) ~ discrete([0.3 : a, 0.7 : d]) ← loan(L) ≅ t.
credit_score(C) ~ gaussian(650, 15.4) ← has_loan(C, L) ≅ Y, Y == f.
credit_score(C) ~ gaussian(700, 10.9) ← has_loan(C, L) ≅ t, status(L) ≅ X, X == a.
credit_score(C) ~ gaussian(600, 20.5) ← has_loan(C, L) ≅ t, status(L) ≅ X, X == d.
    
```

Now, we provide the semantics of such a program.

4.2 Semantics

First, we specify the form of DCs allowed in our programs.

Definition 5. A DC# program is a finite set of DCs $A_0 \sim D \leftarrow A_1, \dots, A_n$ whose atoms A_i ($i > 0$) are either of the following form²:

1. $q(\tau_1, \dots, \tau_k) \cong V$, where $q(\tau_1, \dots, \tau_k)$ is a RV term and V can either be a variable or a constant belonging to the domain of the RV. If V is a variable then it must not appear in RV terms of the DC.
2. $V_1 \diamond V_2$, where V_1, V_2 can be logical variables or constants belonging to the corresponding domains of RVs, and \diamond is a comparison infix predicate that can be either of these: $=, <, >, \geq, \leq$. The predicates \diamond have the same meaning as they have in Prolog.

For example, all clauses shown in Section 4.1 are already of this form, and the program in Example 7 is a DC# program. The reason why logical variables V_i , in the above-stated form (Condition 1), are not allowed to appear in RV terms is that the existence of RVs is not uncertain in DC# programs. This is not the case in the following program:

```
loan_id ~ poisson(10).
status(L) ~ discrete([0.3 : a, 0.7 : d]) ← loan_id ≅ L.
```

Here, the first clause models the identity of loans as a Poisson distribution with a mean of 10. Thus, `loan_id` can be any natural number starting with 0, and the existence of the statuses of loans, say `status(12)`, is uncertain. We disallow writing such open-universe models in the DC# framework because identifying RVs in such models may require analyzing the programs dynamically. It is not clear how to exploit CSIs in such a case. For closed-universes, they can be identified using simple static analysis, which we describe next.

Definition 6. Let \mathbb{P} be a DC# program. An RV set of the program, denoted $rv(\mathbb{P})$, is the set of definite clauses obtained by transforming each clause $A_0 \sim D \leftarrow A_1, \dots, A_n \in \mathbb{P}$ as follows:

1. Let `Body` be the empty set.
2. For each atom A_i ($i > 0$) of the form $q(\tau_1, \dots, \tau_k) \cong V$, an atom $rv(q(\tau_1, \dots, \tau_k))$ is added to `Body`.
3. A clause $rv(A_0) \leftarrow \text{Body}$ is added to $rv(\mathbb{P})$.

Notice that we ignore comparison atoms while constructing the RV set since they do not contain RV terms. They only deal with the values of RVs.

2. Even though only two forms of atoms are allowed, one can still write a deterministic atom $u(\tau_1, \dots, \tau_n)$ as in definite clauses like this: $u(\tau_1, \dots, \tau_n) \cong \text{true}$. So, by restricting the form of allowed atoms, we are not losing expressivity. Recall that a definite clause can be expressed as a DC.

Example 8. *The RV set for the program in Example 7 is:*

```

rv(client(ann)).
rv(loan(l_1)).
rv(loan(l_2)).
rv(has_loan(C,L)) ← rv(client(C),rv(loan(L))).
rv(status(L)) ← rv(loan(L)).
rv(credit_score(C)) ← rv(has_loan(C,L)).
rv(credit_score(C)) ← rv(has_loan(C,L),rv(status(L))).

```

Atom $rv(\text{has_loan}(\text{ann}, l_1))$ is in the least Herbrand model of the above RV set but $rv(\text{has_loan}(l_1, l_2))$ is not. So, term $\text{has_loan}(\text{ann}, l_1)$ is identified as a RV but the meaningless term $\text{has_loan}(l_1, l_2)$ is not, even though it belongs to the Herbrand universe. Recall that a ground instance of a DC defines a RV only when its body is true. No ground DC defines $\text{has_loan}(l_1, l_2)$ as a RV since term l_1 is a loan and not a client. We will show that the RV set of a program identifies all RVs defined by the program. This is similar to Bayesian clauses in Bayesian logic programs, where atoms in the least Herbrand model of Bayesian clauses are RVs over which a probability distribution is defined (Kersting & De Raedt, 2007).

Now, it is possible to ground a DC# program \mathbb{P} given an assignment \mathbf{u} of RVs. We denote such ground programs by $\text{ground}(\mathbb{P})_{\mathbf{u}}$.

Example 9. *Given the following assignment of RVs identified from the RV set (Example 8),*

$$\mathbf{u} = \{\text{client}(\text{ann}) \cong t, \text{loan}(l_1) \cong t, \text{loan}(l_2) \cong t, \text{has_loan}(\text{ann}, l_1) \cong t, \\ \text{has_loan}(\text{ann}, l_2) \cong t, \text{status}(l_1) \cong a, \text{status}(l_2) \cong d, \text{credit_score}(\text{ann}) \cong 601.2\},$$

the program of Example 7 grounds with respect to the assignments like this:

```

client(ann) ~ val(t).
loan(l_1) ~ val(t).
loan(l_2) ~ val(t).
has_loan(ann, l_1) ~ bernoulli(0.2) ← client(ann) ≅ t, loan(l_1) ≅ t.
has_loan(ann, l_2) ~ bernoulli(0.2) ← client(ann) ≅ t, loan(l_2) ≅ t.
status(l_1) ~ discrete([0.3 : a, 0.7 : d]) ← loan(l_1) ≅ t.
status(l_2) ~ discrete([0.3 : a, 0.7 : d]) ← loan(l_2) ≅ t.
credit_score(ann) ~ gaussian(650, 15.4) ← has_loan(ann, l_1) ≅ t, t == f.
credit_score(ann) ~ gaussian(650, 15.4) ← has_loan(ann, l_2) ≅ t, t == f.
credit_score(ann) ~ gaussian(700, 10.9) ← has_loan(ann, l_1) ≅ t, status(l_1) ≅ a, a == a.
credit_score(ann) ~ gaussian(700, 10.9) ← has_loan(ann, l_2) ≅ t, status(l_2) ≅ d, d == a.
credit_score(ann) ~ gaussian(600, 20.5) ← has_loan(ann, l_1) ≅ t, status(l_1) ≅ a, a == d.
credit_score(ann) ~ gaussian(600, 20.5) ← has_loan(ann, l_2) ≅ t, status(l_2) ≅ d, d == d.

```

Notice that neither meaningless terms like $\text{has_loan}(l_1, l_2)$ appear in this ground program nor atoms like $\text{has_loan}(\text{ann}, l_1) \cong f$ appear.

One might wonder why we require RVs and their assignment to be *given* before grounding DC# programs. There are two main reasons. First, the programs may define continuous RVs that can take infinitely many values, so we would be constructing infinitely large ground programs without knowing the values of RVs. Second, we do not want meaningless terms like `has_loan(1.1,1.2)` to appear in the head of clauses in the ground programs. These clauses define RVs, and it does not make sense to treat these meaningless terms as RVs. So, we should know RVs before grounding the programs.

Furthermore, the assignment \mathbf{u} can be thought of as asserted facts in the ground programs, which decide the truth values of atoms of the form $X \cong Y$ in the body of clauses. Since the truth values of bodies of clauses depend on the assignment \mathbf{u} of RV terms, when the body of a clause is true, we say that the body is true *with respect to* \mathbf{u} .

Clearly, the next step is to identify direct influence relationships among RVs defined by programs.

Definition 7. *Let \mathbb{P} be a DC# program, \mathbf{u} be an assignment of RVs, and $\text{ground}(\mathbb{P})_{\mathbf{u}}$ be a ground program constructed given \mathbf{u} . A ground RV term A directly influences B if there is clause $B \sim D \leftarrow B_1, \dots, B_n \in \text{ground}(\mathbb{P})_{\mathbf{u}}$ for some \mathbf{u} such that $n > 0$, A is in the body of the clause, and the body is true with respect to \mathbf{u} .*

For example, we observe from the ground program of Example 9 that ground RV terms `has_loan(ann,1.1)`, `has_loan(ann,1.2)`, `status(1.1)`, and `status(1.2)` directly influence `credit_score(ann)`. However, these relationships are defined with respect to assignments \mathbf{u} , and it is hard to construct ground programs with respect to all \mathbf{u} for identifying these relationships. For this purpose, we need a simple approach that can be implemented and executed efficiently. Next, we present such an approach. It performs a simple transformation of the RV sets of DC# programs.

Definition 8. *Let $\text{rv}(\mathbb{P})$ be the RV set of a DC# program \mathbb{P} . The dependency set $\text{dep}(\mathbb{P})$ is the union of $\text{rv}(\mathbb{P})$ and the set of definite clauses $\text{pa}(\mathbb{P})$ obtained by transforming each clause $\text{rv}(A_0) \leftarrow \text{rv}(A_1), \dots, \text{rv}(A_n) \in \text{rv}(\mathbb{P})$ with non empty body as follows:*

- *For each $\text{rv}(A_i)$, such that $i > 0$, a clause $\text{pa}(A_0, A_i) \leftarrow \text{rv}(A_0), \text{rv}(A_1), \dots, \text{rv}(A_n)$ is added to $\text{pa}(\mathbb{P})$.*

Example 10. *By transforming the RV set constructed in Example 8, we obtain the dependency set of Example 7 that consists of the RV set and the following extra clauses:*

```

pa(has_loan(C,L), client(C)) ← rv(has_loan(C,L)), rv(client(C)), rv(loan(L)).
pa(has_loan(C,L), loan(L)) ← rv(has_loan(C,L)), rv(client(C)), rv(loan(L)).
pa(status(L), loan(L)) ← rv(status(L)), rv(loan(L)).
pa(credit_score(C), has_loan(C,L)) ← rv(credit_score(C)), rv(has_loan(C,L)).
pa(credit_score(C), has_loan(C,L)) ← rv(credit_score(C)), rv(has_loan(C,L)), rv(status(L)).
pa(credit_score(C), status(L)) ← rv(credit_score(C)), rv(has_loan(C,L)), rv(status(L)).

```

One can easily infer from the above definite program that `status(1.1)` directly influences `credit_score(ann)` since $\text{pa}(\text{credit_score}(\text{ann}), \text{status}(1.1))$ is entailed by the above program. Basically, $\text{pa}(A, B)$ states that the parent of A is B . Note that Definition 7 defines the direct influences and Definition 8 presents an approach to identify them efficiently. The

dependency set is similar to the dependency graph in BLPs (Kersting & De Raedt, 2007). The only difference is that we use definite clauses instead of graphs to represent direct influences. These clauses can be automatically constructed from DC# programs using the transformations.

However, it is still unclear how to interpret the clauses, in the ground programs, which may specify multiple distributions for RVs.

Definition 9. Let $\text{ground}(\mathbb{P})_{\mathbf{u}}$ be a ground DC# program constructed given an assignment \mathbf{u} , and let

$$\begin{aligned} A_0 \sim D_{A_{10}} &\leftarrow A_{11}, \dots, A_{1n_1} \\ &\vdots \\ A_0 \sim D_{A_{k0}} &\leftarrow A_{k1}, \dots, A_{kn_k} \end{aligned}$$

be k clauses for A_0 in $\text{ground}(\mathbb{P})_{\mathbf{u}}$, whose bodies are true with respect to \mathbf{u} . Then, there is a multiset of distributions $[D_{A_{10}}, \dots, D_{A_{k0}}]$ specified for A_0 in $\text{ground}(\mathbb{P})_{\mathbf{u}}$. If $k > 1$ for some A_0 then we say the clauses in $\text{ground}(\mathbb{P})_{\mathbf{u}}$ are mutually inclusive; otherwise, they are mutually exclusive.

For example, there are two distributions specified for `credit_score(ann)` in the above example since client `ann` has two loans and these loans make the bodies of two clauses for `credit_score(ann)` true according to the used assignment of RVs. This, however, raises several questions: What distribution does the RV follow? Using DCs, how can we describe that multiple loans probabilistically influence the credit score?

The standard answer to these questions is to combine multiple distributions into a single distribution using so-called *combining rules* (Ngo & Haddawy, 1995; Kersting & De Raedt, 2001; Jaeger, 2007; Natarajan et al., 2008). Combining rules are based on the assumption of *independence of causal influence* (ICI), where it is assumed that multiple causes on a target variable can be decomposed into several independent causes whose effects are combined to yield a final value (Zhang & Poole, 1996). In terms of the above example, this means that we can let each loan independently define a distribution for the credit score and then somehow combine all the defined distributions into a single distribution using a combining rule \mathcal{CR} . The most commonly used combining rules in relational systems are Mean, and NoisyOR (Kersting & De Raedt, 2007; Natarajan et al., 2008; Fierens et al., 2015).

Let $[\mathcal{D}_1, \dots, \mathcal{D}_n]$ be a non-empty multiset of probability distributions. Then, the mean combining rule is defined as follows:

$$\text{Mean}([\mathcal{D}_1, \dots, \mathcal{D}_n]) = \frac{1}{n} \sum_{i=1}^n \mathcal{D}_i$$

where the right-hand side of the equation is a mixture of distributions. When all distributions in the multiset are Bernoulli distributions, denoted `bernoulli(pi)`, with parameters p_i , then generally noisy OR combining rule is used, which is defined as follows:

$$\text{NoisyOR}([\text{bernoulli}(p_1), \dots, \text{bernoulli}(p_n)]) = \text{bernoulli}(1 - \prod_{i=1}^n (1 - p_i))$$

To specify a proper probability distribution, a DC# program has to be *well-defined*.

Definition 10. Let \mathbb{P} be a DC# program, $\mathbf{V}_{\mathbb{P}}$ be the set of all RVs identified from $rv(\mathbb{P})$, ω be an assignment of $\mathbf{V}_{\mathbb{P}}$, and Ω be the set of all such ω . The program \mathbb{P} is well-defined if it satisfies the following conditions:

1. Exhaustiveness: For all $\mathbf{A}_0 \in \mathbf{V}_{\mathbb{P}}$ and for all $\omega \in \Omega$, there is at least one clause of the form $\mathbf{A}_0 \sim \mathbf{D} \leftarrow \mathbf{A}_1, \dots, \mathbf{A}_n$ in $\text{ground}(\mathbb{P})_{\omega}$ such that the clause's body is true with respect to ω .
2. Acyclicity: There exists a function $\text{rank}(\cdot)$ that maps $\mathbf{A} \in \mathbf{V}_{\mathbb{P}}$ to natural numbers \mathbb{N} . Let $\mathbf{A}_0 \sim \mathbf{D} \leftarrow \mathbf{A}_1, \dots, \mathbf{A}_n$ be a clause in $\text{ground}(\mathbb{P})_{\omega}$, and let $\{\mathbf{rv}_1, \dots, \mathbf{rv}_m\}$ be the set of RV terms in $\{\mathbf{A}_1, \dots, \mathbf{A}_n\}$. For all ω , for all clauses in $\text{ground}(\mathbb{P})_{\omega}$, and for all i , $\text{rank}(\mathbf{A}_0) > \text{rank}(\mathbf{rv}_i)$.
3. Finite Support: The set $\mathbf{V}_{\mathbb{P}}$ is non-empty and each $\mathbf{A}_0 \in \mathbf{V}_{\mathbb{P}}$ is directly influenced by a finite set of RVs.

These conditions are similar to the conditions imposed on BLPs for them to be well-defined (Kersting & De Raedt, 2007). Since BLPs consist of Bayesian clauses along with CPDs, the exhaustiveness condition is implicitly imposed. Recall that CPDs define a conditional probability distribution for an RV given any possible assignment of the RV's parents, which means CPDs are exhaustive.

Let us investigate some *ill-defined* programs.

Example 11. The program

$$\mathbf{a}(\mathbf{X}) \sim \text{bernoulli}(0.2) \leftarrow \mathbf{b}(\mathbf{X}) \cong \mathbf{t}.$$

is not well-defined since the least Herbrand model of the RV set of the program is empty. The program

$$\begin{aligned} \mathbf{s}(\mathbf{a}, \mathbf{b}) &\sim \text{val}(\mathbf{t}). \\ \mathbf{s}(\mathbf{X}, \mathbf{f}(\mathbf{Y})) &\sim \text{val}(\mathbf{t}) \leftarrow \mathbf{s}(\mathbf{X}, \mathbf{Y}) \cong \mathbf{t}. \\ \mathbf{r}(\mathbf{X}) &\sim \text{val}(\mathbf{t}) \leftarrow \mathbf{s}(\mathbf{X}, \mathbf{f}(\mathbf{Y})) \cong \mathbf{t}. \end{aligned}$$

is ill-defined because $\mathbf{r}(\mathbf{a})$ is influenced by infinite number of RVs: $\mathbf{s}(\mathbf{a}, \mathbf{b})$, $\mathbf{s}(\mathbf{a}, \mathbf{f}(\mathbf{b}))$, $\mathbf{s}(\mathbf{a}, \mathbf{f}(\mathbf{f}(\mathbf{b})))$, and so on. However, the following program

$$\begin{aligned} \mathbf{s}(\mathbf{a}, \mathbf{b}) &\sim \text{val}(\mathbf{t}). \\ \mathbf{s}(\mathbf{X}, \mathbf{f}(\mathbf{Y})) &\sim \text{val}(\mathbf{t}) \leftarrow \mathbf{s}(\mathbf{X}, \mathbf{Y}) \cong \mathbf{t}. \end{aligned}$$

is well-defined since each RV is influenced by a finite number of RVs. Such programs allow for writing models that can have infinite RVs such as hidden Markov models. The following program violates the exhaustiveness condition.

$$\begin{aligned} \mathbf{a}(1) &\sim \text{discrete}([0.2 : \mathbf{t}, 0.8 : \mathbf{f}]). \\ \mathbf{b}(1) &\sim \text{bernoulli}(0.6) \leftarrow \mathbf{a}(1) \cong \mathbf{t}. \end{aligned}$$

This is because the distribution of $\mathbf{b}(1)$ when $\mathbf{a}(1)$ is \mathbf{f} (“false”) is undefined. The program

$$\begin{aligned} \mathbf{a}(1) &\sim \text{discrete}([0.2 : \mathbf{t}, 0.8 : \mathbf{f}]). \\ \mathbf{a}(1) &\sim \text{discrete}([0.1 : \mathbf{t}, 0.9 : \mathbf{f}]) \leftarrow \mathbf{a}(1) \cong \mathbf{t}. \\ \mathbf{a}(1) &\sim \text{discrete}([0.7 : \mathbf{t}, 0.3 : \mathbf{f}]) \leftarrow \mathbf{a}(1) \cong \mathbf{f}. \end{aligned}$$

is ill-defined as well since it has a cyclic dependency.

We can show that the RV set of a well-defined DC# program identifies all RVs defined by the program, and the dependency set all direct influences among RVs.

Proposition 2. *Let $rv(\mathbb{P})$ be a RV set of a well-defined DC# program \mathbb{P} . Then, \mathbb{P} defines \mathbf{A} as a RV iff $rv(\mathbf{A})$ is in the least Herbrand model of $rv(\mathbb{P})$.*

The proofs of all results presented in this paper are in Appendix A.

Proposition 3. *Let $dep(\mathbb{P})$ be a dependency set of a well-defined DC# program \mathbb{P} . Then, \mathbf{A} directly influences \mathbf{B} iff $pa(\mathbf{B}, \mathbf{A})$ is in the least Herbrand model of $dep(\mathbb{P})$.*

A *possible world* is an assignment of all RVs $\mathbf{V}_{\mathbb{P}}$ defined by a program. Since $\mathbf{V}_{\mathbb{P}}$ can be an infinite set, as discussed in Example 11, defining how probabilities are assigned to possible worlds is non-trivial³. However, we can still define how probabilities are assigned to assignments of a certain subset of RVs.

Definition 11. *Let $\mathbf{V}_{\mathbb{P}}$ be the set of all RVs defined by a well-defined DC# program \mathbb{P} , \mathbf{u} be an assignment of a finite subset $\mathbf{U} \subseteq \mathbf{V}_{\mathbb{P}}$. Denote the set of RVs that directly influence X by $\mathbf{Pa}(X)$. The assignment \mathbf{u} is said to be closed under direct influence relationships if $X \in \mathbf{U}$ implies $\mathbf{Pa}(X) \subseteq \mathbf{U}$.*

Example 12. *Reconsider the DC# program of Example 7. Partial assignment $\mathbf{u}_1 = \{\text{client}(\text{ann}) \cong \mathbf{t}, \text{loan}(\mathbf{l}, \mathbf{l}) \cong \mathbf{t}, \text{has_loan}(\text{ann}, \mathbf{l}, \mathbf{l}) \cong \mathbf{f}\}$ is closed under direct influence relationships, but partial assignment $\mathbf{u}_2 = \{\text{credit_score}(\text{ann}) \cong 651.2, \text{client}(\text{ann}) \cong \mathbf{t}, \text{loan}(\mathbf{l}, \mathbf{l}) \cong \mathbf{t}, \text{has_loan}(\text{ann}, \mathbf{l}, \mathbf{l}) \cong \mathbf{f}\}$ is not closed under such relationships. This is because random variable term $\text{has_loan}(\text{ann}, \mathbf{l}, \mathbf{l}, \mathbf{l})$ directly influences $\text{credit_score}(\text{ann})$, but is not assigned in \mathbf{u}_2 .*

Definition 12. *Let \mathbb{P} be a well-defined DC# program, \mathbf{u} be an assignment that is closed under direct influence relationships, and $\text{ground}(\mathbb{P})_{\mathbf{u}}$ be the ground program constructed given \mathbf{u} . Then the probability (or density) that $\text{ground}(\mathbb{P})_{\mathbf{u}}$ assigns to \mathbf{u} is given by*

$$P(\mathbf{u}) = \prod_{\mathbf{A}_i \cong \mathbf{x} \in \mathbf{u}} \text{CR}([\mathbf{D}_{\mathbf{A}_{i1}}, \dots, \mathbf{D}_{\mathbf{A}_{ki}}])(\mathbf{A}_i \cong \mathbf{x})$$

where $[\mathbf{D}_{\mathbf{A}_{i1}}, \dots, \mathbf{D}_{\mathbf{A}_{ki}}]$ is a multiset of distributions specified for RV term \mathbf{A}_i in $\text{ground}(\mathbb{P})_{\mathbf{u}}$ and $\text{CR}([\mathbf{D}_{\mathbf{A}_{i1}}, \dots, \mathbf{D}_{\mathbf{A}_{ki}}])(\mathbf{A}_i \cong \mathbf{x})$ is the conditional probability density/mass of \mathbf{A}_i at \mathbf{x} according to the combined distribution $\text{CR}([\mathbf{D}_{\mathbf{A}_{i1}}, \dots, \mathbf{D}_{\mathbf{A}_{ki}}])$ obtained after applying CR on the multiset.

3. To define a probability distribution over infinite RVs, one uses the Kolmogorov extension theorem.

Example 13. Suppose we use Mean as a combining rule for the program of Example 9. Then, to compute the probability density given to the assignment in the example, we will use the following conditional probability densities/masses:

<code>client(ann) ≅ t</code>	1.00
<code>loan(1_1) ≅ t</code>	1.00
<code>loan(1_2) ≅ t</code>	1.00
<code>has_loan(ann, 1_1) ≅ t</code>	0.20
<code>has_loan(ann, 1_2) ≅ t</code>	0.20
<code>status(1_1) ≅ a</code>	0.30
<code>status(1_2) ≅ d</code>	0.70
<code>credit_score(ann) ≅ 601.2</code>	0.04

where, the last entry is the probability density at 601.2 according to the mixture of Gaussians $\{\mathcal{N}(x \mid 700, 10.9) + \mathcal{N}(x \mid 600, 20.5)\}/2$. Thus, the density assigned to the assignment is $0.20 \times 0.20 \times 0.30 \times 0.70 \times 0.04$.

We can show that these probability (or density) assignments define a unique probability distribution.

Proposition 4. Let \mathbb{P} be a well-defined DC# program, and $\mathbf{V}_{\mathbb{P}}$ be a set of all RVs defined by it. Then \mathbb{P} specifies a unique probability distribution over $\mathbf{V}_{\mathbb{P}}$.

Since well-defined DC# programs satisfy all conditions of well-defined BLPs, this proposition follows from Proposition 1 of (Kersting & De Raedt, 2001)

5. Inference in Ground DC# Programs

Before presenting the sampling algorithm for first-order DC# programs, we will first design an efficient sampling algorithm for ground DC# programs describing BNs with structured CPDs. Thus, this section will focus only on programs having mutually exclusive clauses, which are sufficient to describe such BNs. The algorithm presented in this section aims to exploit structures within CPDs or the structure of clauses in ground programs. However, in this section, we will assume that continuous RVs are absent for simplicity. Once this algorithm is clear, the extended algorithm for first-order DC# programs presented in the next section will be easy to comprehend.

A natural representation of the structures in CPDs is via *tree-CPDs* (Koller & Friedman, 2009), as illustrated in Figure 1a. For all assignments to the parents of a variable A , a unique leaf in the tree specifies a (conditional) distribution over A . The path to each leaf dictates the contexts, i.e., *partially assigned parents*, given which this distribution is used. We can easily represent tree-CPDs using DCs, where each path from the root to a leaf in each tree-CPD maps to a rule.

Example 14. The set of clauses for the tree-CPD in Figure 1a:

- $e \sim \text{bernoulli}(0.2) \leftarrow a \cong 1.$
- $e \sim \text{bernoulli}(0.9) \leftarrow a \cong 0, b \cong 1.$
- $e \sim \text{bernoulli}(0.6) \leftarrow a \cong 0, b \cong 0, c \cong 1.$
- $e \sim \text{bernoulli}(0.3) \leftarrow a \cong 0, b \cong 0, c \cong 0.$

Conversely, we can view a ground program with mutually exclusive clauses as a representation of tree-CPDs of all variables of a BN. Thus, we can alternatively define the probability distribution specified by such ground programs as follows,

Definition 13. *Let \mathcal{B} be a Bayesian network with tree-CPDs specifying a distribution P . Let \mathbb{P} be a set of distributional clauses such that each path from the root to a leaf of each tree-CPD corresponds to a clause in \mathbb{P} . Then \mathbb{P} specifies the same distribution P .*

To avoid confusion, such ground programs will be called $\text{DC}(\mathcal{B})$ programs.

While an efficient sampling algorithm for \mathcal{B} only exploits the graph structure (CIs properties) in \mathcal{B} , the key to designing an efficient sampling algorithm for $\text{DC}(\mathcal{B})$ programs is to exploit both the underlying graph and the clause structure (CSIs properties). To this end, we start with our discussion on the estimation of unconditional probability queries, which is necessary to support our further discussion on conditional probability queries, where we present the full algorithm for $\text{DC}(\mathcal{B})$ programs.

5.1 Top-Down Proof Procedure for $\text{DC}(\mathcal{B})$ programs

Estimating unconditional probability queries in BNs is easy. We just need to generate some random samples of query variables and to find the fraction of times the query is true, which is the estimated probability of the query. In this sampling process, all ancestors of query variables are also sampled. However, due to the clause structure in $\text{DC}(\mathcal{B})$ programs, it is possible to generate random samples of query variables by sampling only some (not all) ancestors, which makes the sampling process more efficient. A simplified version of an approach due to (Nitti et al., 2016) is discussed in Algorithm 1. This approach resembles SLD resolution (Kowalski, 1974) for definite programs. However, there are some differences due to the stochastic nature of sampling. Unlike SLD resolution, this approach maintains global variable \mathbf{Asg} to record sampled values of RVs.

Given an initial goal G_0 and global variable \mathbf{Asg} , the algorithm recursively produces new goals G_1, G_2, \dots , and updates \mathbf{Asg} . There are two cases when it is not possible to obtain G_{i+1} from G_i :

- the first is when the selected subgoal of the form $t \cong v$ cannot be resolved because the sampled value of the RV term already recorded in \mathbf{Asg} is different from v .
- the other case appears when $G_i = \square$ (i.e. the empty goal).

The procedure⁴ results in a *derivation* of G_0 , a finite sequence of goals starting with the initial goal

Example 15. *Consider the initial goal $(G_0) \leftarrow e \cong 1$ and the following $\text{DC}(\mathcal{B})$ program:*

```

a ~ bernoulli(0.1).
d ~ bernoulli(0.3).
b ~ bernoulli(0.2) ← a ≅ 0.
b ~ bernoulli(0.6) ← a ≅ 1.
c ~ bernoulli(0.2) ← a ≅ 1.

```

4. Since cyclic dependency among RVs is not allowed in well-defined programs, and the number of clauses in the program is finite, the procedure is guaranteed to terminate.

Algorithm 1 DC(\mathcal{B}) Proof Procedure

procedure PROVE-GROUND(\mathbf{G})

- Proves a conjunction of ground atoms \mathbf{G} .
 - Returns **yes** if there is a choice that makes \mathbf{G} empty; otherwise the procedure fails.
1. While \mathbf{G} is not empty:
 - (a) Select the first atom \mathbf{A} from \mathbf{G} .
 - (b) If \mathbf{A} is of the form $\mathbf{t} \cong \mathbf{v}$:
 - i. If a value of \mathbf{t} is recorded in \mathbf{Asg} :
 - A. If $\mathbf{Asg}[\mathbf{t}] == \mathbf{v}$: remove \mathbf{A} from \mathbf{G} .
 - ii. Else:
 - A. **Choose** $\mathbf{t} \sim D \leftarrow B_1, \dots, B_n \in \mathbb{P}$ with \mathbf{t} in the head.
 - B. Set $\mathbf{G} := B_1, \dots, B_n, \mathbf{sample}(\mathbf{t}, D), \mathbf{G}$.
 - (c) Else If \mathbf{A} is of the form $\mathbf{sample}(\mathbf{t}, D)$:
 - i. Sample a value \mathbf{v} from D , record $\mathbf{Asg}[\mathbf{t}] := \mathbf{v}$, and remove \mathbf{A} from \mathbf{G} .
 2. Return **yes**.
-

 $c \sim \text{bernoulli}(0.7) \leftarrow a \cong 0, b \cong 1.$ $c \sim \text{bernoulli}(0.8) \leftarrow a \cong 0, b \cong 0.$ $e \sim \text{bernoulli}(0.9) \leftarrow c \cong 1.$ $e \sim \text{bernoulli}(0.4) \leftarrow c \cong 0, d \cong 1.$ $e \sim \text{bernoulli}(0.3) \leftarrow c \cong 0, d \cong 0.$

Use $\mathbf{sample}(k, 1)$ as a shorthand notation for $\mathbf{sample}(k, \text{bernoulli}(1))$. A derivation of \mathbf{G}_0 , the state of \mathbf{Asg} and the program clause used in each step is shown in Figure 3.

As usual, the derivation of goal \mathbf{G}_0 that ends in the empty goal corresponds to a *refutation* of the goal. Not all derivations lead to refutations. As already pointed out, if the selected subgoal cannot be resolved, the derivation *fails*. A refutation or a failed derivation is a *complete derivation*. If the selected subgoal of some goal can be resolved with more than one program clause, there can be many complete derivations. The procedure searches a refutation by generating multiple complete derivations in a depth-first search fashion. Notice that, unlike SLD-resolution, here derivations might depend on previous complete derivations. This is because the previous derivations might have updated the \mathbf{Asg} table, which in turn might influence the current derivation. Additionally, it should be clear that no or only one derivation can be the refutation when clauses in a program are mutually exclusive. So, for efficiency, one should stop generating more derivations in such programs once a refutation is obtained.

To estimate the probability of $\mathbf{e} \cong 1$, we call PROVE-GROUND($\mathbf{e} \cong 1$) repeatedly. The fraction of times we get refutation (i.e., the algorithm returns **yes**) is the estimated probability. Notice that some ancestors of the query variables may not be sampled on some occasions, e.g., in Example 15, variables \mathbf{d} and \mathbf{b} were not sampled. Thus, we speed up the

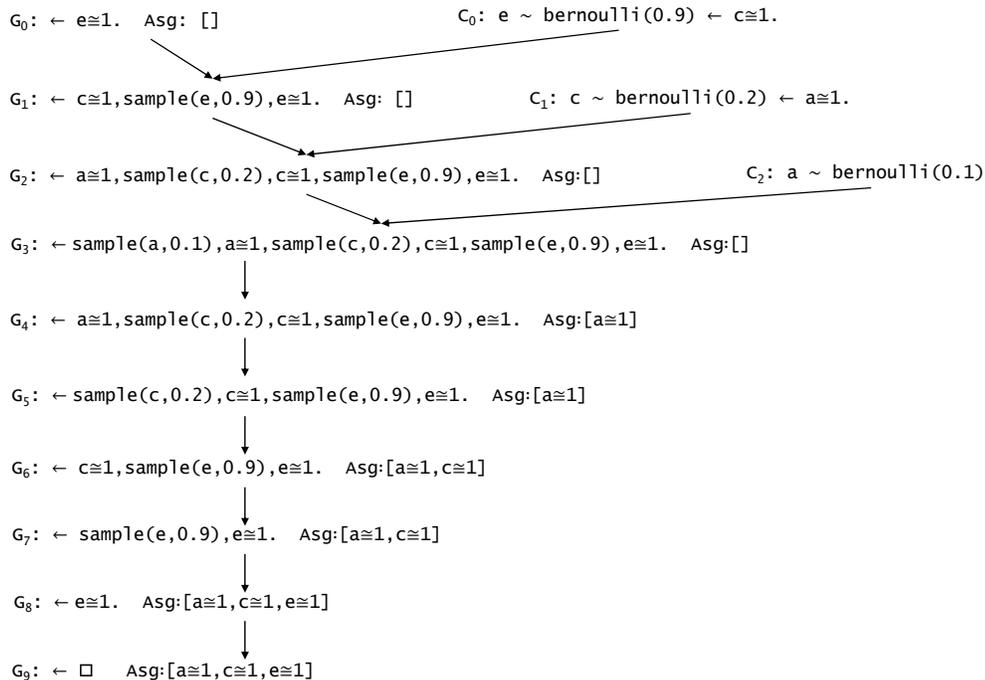


Figure 3: A search graph for a top-down derivation

sampling process and, in this way, exploit the structure of clauses while estimating marginal probabilities.

5.2 Exploiting Context-Specific Independencies

We now turn our attention to the problem of estimating conditional probabilities, which is more interesting but is significantly more complicated. Exploiting CSIs in this problem is unexplored in (Nitti et al., 2016). Here we introduce a sampling algorithm that combines the proof procedure, discussed in the previous section, with the Bayes-ball simulation, discussed in Section 3.2. The combined algorithm then exploits both the clause structures and the underlying graph structure of $DC(\mathcal{B})$ programs. As we will see, it samples variables given the states of only some of their requisite ancestors. This contrasts with the Bayes-ball simulation of BNs, where knowledge of all such ancestors’ states is required. This section is divided into two parts. The first part presents a novel notion of contextual assignment that allows for exploiting CSIs. It provides insight into the computation of μ using partial assignments of requisite variables. We will show that CSIs allow for breaking the main problem of computing μ into several sub-problems that can be solved independently. The second part presents the sampling algorithm and justifies it using the notion introduced in the first part.

5.2.1 NOTION OF CONTEXTUAL ASSIGNMENTS

Recall from Section 3.2 that variables \mathbf{X} , \mathbf{Z}_\star , \mathbf{e}_\star , \mathbf{E}_\star are requisite for computing the query μ and the requisite network \mathcal{B}_\star is formed by these variables. We will consider partial assignments of these variables, which will be used to compute μ . Let us start by defining these assignments.

Definition 14. Let $\mathbf{Z}_\dagger \subseteq \mathbf{Z}_\star$ and $\mathbf{e}_\dagger \subseteq \mathbf{e}_\star$. Denote $\mathbf{Z}_\star \setminus \mathbf{Z}_\dagger$ by \mathbf{Z}_{\ddagger} , and $\mathbf{e}_\star \setminus \mathbf{e}_\dagger$ by \mathbf{e}_{\ddagger} . A partial assignment \mathbf{x} , \mathbf{z}_\dagger , \mathbf{Z}_{\ddagger} , \mathbf{e}_\dagger , \mathbf{e}_{\ddagger} will be called contextual assignment if due to CSIs in P ,

$$\prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(u_i \mid \mathbf{pa}(U_i)) = \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(u_i \mid \mathbf{ppa}(U_i))$$

where $\mathbf{ppa}(U_i)$ is a set of partially assigned parents of U_i , that is, $\mathbf{Ppa}(U_i) = \mathbf{Pa}(U_i) \setminus \mathbf{Z}_{\ddagger}$.

Here, we summarise some important RVs or their assignments:

\mathbf{X}	Query variables
\mathbf{Z}_\star	Unobserved requisite variables apart from query variables
\mathbf{e}_\star	Diagnostic evidence
\mathbf{e}_\star	Predictive evidence
\mathbf{z}_\dagger	Assigned subset of \mathbf{Z}_\star in a contextual assignment
\mathbf{Z}_{\ddagger}	Subset of \mathbf{Z}_\star not assigned in a contextual assignment
\mathbf{e}_\dagger	Subset of \mathbf{e}_\star in a contextual assignment
\mathbf{e}_{\ddagger}	Subset of \mathbf{e}_\star not in a contextual assignment

Example 16. Consider the network of Figure 1a, and assume that our diagnostic evidence is $\{F = 1, G = 0, H = 1\}$, predictive evidence is $\{D = 1\}$, and query is $\{E = 0\}$. From the CPD's structure, we have: $P(E = 0 \mid A = 1, B, C) = P(E = 0 \mid A = 1)$; consequently, a contextual assignment is $\mathbf{x} = \{E = 0\}$, $\mathbf{z}_\dagger = \{A = 1\}$, $\mathbf{e}_\dagger = \{\}$, $\mathbf{Z}_{\ddagger} = \{B, C\}$, $\mathbf{e}_{\ddagger} = \{F = 1, G = 0, H = 1\}$. We also have: $P(E = 0 \mid A = 0, B = 1, C) = P(E = 0 \mid A = 0, B = 1)$; consequently, another such assignment is $\mathbf{x} = \{E = 0\}$, $\mathbf{z}_\dagger = \{A = 0, B = 1\}$, $\mathbf{e}_\dagger = \{H = 1\}$, $\mathbf{Z}_{\ddagger} = \{C\}$, $\mathbf{e}_{\ddagger} = \{F = 1, G = 0\}$.

We aim to treat the evidence \mathbf{e}_{\ddagger} independently, thus, we define it first.

Definition 15. The diagnostic evidence \mathbf{e}_{\ddagger} in a contextual assignment \mathbf{x} , \mathbf{z}_\dagger , \mathbf{Z}_{\ddagger} , \mathbf{e}_\dagger , \mathbf{e}_{\ddagger} will be called residual evidence.

However, contextual assignments do not immediately allow us to treat the residual evidence independently. We need the assignments to be safe.

Definition 16. Let $e \in \mathbf{e}_\star$ be a diagnostic evidence, and let S be an unobserved ancestor of E in the graph structure in \mathcal{B}_\star , where \mathcal{B}_\star is the sub-network formed by the requisite variables. Let $S \rightarrow \dots B_i \dots \rightarrow E$ be a causal trail such that either no B_i is observed or there is no B_i . Let \mathbf{S} be the set of all such S . Then the variables \mathbf{S} will be called basis of e . Let $\dot{\mathbf{e}}_\star \subseteq \mathbf{e}_\star$, and let $\dot{\mathbf{S}}_\star$ be the set of all such S for all $e \in \dot{\mathbf{e}}_\star$. Then $\dot{\mathbf{S}}_\star$ will be called basis of $\dot{\mathbf{e}}_\star$.

Reconsider Example 16; the basis of $\{F = 1\}$ is $\{B\}$.

Definition 17. Let \mathbf{x} , \mathbf{z}_\dagger , \mathbf{Z}_\dagger , \mathbf{e}_\dagger , \mathbf{e}_\ddagger be a contextual assignment, and let \mathbf{S}_\dagger be the basis of the residual evidence \mathbf{e}_\dagger . If $\mathbf{S}_\dagger \subseteq \mathbf{Z}_\dagger$ then the contextual assignment will be called safe.

Example 17. Reconsider Example 16; the first example of a contextual assignment is safe, but the second is not since the basis B of \mathbf{e}_\dagger has a non-empty intersection with \mathbf{Z}_\dagger . We can make the second safe like this: $\mathbf{x} = \{E = 0\}$, $\mathbf{z}_\dagger = \{A = 0, B = 1\}$, $\mathbf{e}_\dagger = \{F = 1, H = 1\}$, $\mathbf{Z}_\dagger = \{C\}$, $\mathbf{e}_\ddagger = \{G = 0\}$. See Figure 4.

Before showing that the residual evidence can now be treated independently, we first define a random variable called *weight*.

Definition 18. Let $e \in \mathbf{e}_\star$ be a diagnostic evidence, and let W_e be a random variable defined as follows:

$$W_e = P(e \mid \mathbf{Pa}(E)).$$

The variable W_e will be called weight of e . The weight of a subset $\dot{\mathbf{e}}_\star \subseteq \mathbf{e}_\star$ is defined as follows:

$$W_{\dot{\mathbf{e}}_\star} = \prod_{u_i \in \dot{\mathbf{e}}_\star} P(u_i \mid \mathbf{Pa}(U_i)).$$

Now we can show the following result:

Theorem 5. Let $\dot{\mathbf{e}}_\star \subseteq \mathbf{e}_\star$, and let $\dot{\mathbf{S}}_\star$ be the basis of $\dot{\mathbf{e}}_\star$. Then the expectation of weight $W_{\dot{\mathbf{e}}_\star}$ relative to the distribution Q_\star as defined in Equation 4 can be written as:

$$\mathbb{E}_{Q_\star}[W_{\dot{\mathbf{e}}_\star}] = \sum_{\dot{\mathbf{S}}_\star} \prod_{u_i \in \dot{\mathbf{e}}_\star \cup \dot{\mathbf{S}}_\star} P(u_i \mid \mathbf{pa}(U_i)).$$

Hence, apart from unobserved variables $\dot{\mathbf{S}}_\star$, the computation of $\mathbb{E}_{Q_\star}[W_{\dot{\mathbf{e}}_\star}]$ does not depend on other unobserved variables.

Let ψ denotes a contextual assignment \mathbf{x} , \mathbf{z}_\dagger , \mathbf{Z}_\dagger , \mathbf{e}_\dagger , \mathbf{e}_\ddagger . The range of ψ , denoted $\text{range}(\psi)$, is the set of all full assignments constructed by assigning the unobserved variables in \mathbf{Z}_\dagger . Reconsider the contextual assignment of Example 17. Since C is a Boolean random variable, the range of the contextual assignment is $\{\{E = 0, A = 0, B = 1, F = 1, H = 1, C = 0, G = 0\}, \{E = 0, A = 0, B = 1, F = 1, H = 1, C = 1, G = 0\}\}$.

It is worth noting that a full assignment ψ is a safe contextual assignment, where the set of unobserved variables and the residual evidence set are empty. In such a case, $\text{range}(\psi) = \{\psi\}$.

The next theorem requires contextual assignments to be mutually exclusive. Two contextual assignments ψ and ψ' are mutually exclusive if $\text{range}(\psi) \cap \text{range}(\psi') = \emptyset$.

Theorem 6. Let Ψ be a set of mutually exclusive contextual assignments such that each full assignment \mathbf{x} , \mathbf{z}_\star , \mathbf{e}_\star , \mathbf{e}_\ddagger of the requisite network \mathcal{B}_\star is under the range of a safe contextual assignment $\psi \in \Psi$. Let $\mathbf{x}[\psi]$, $\mathbf{z}_\dagger[\psi]$, $\mathbf{Z}_\dagger[\psi]$, $\mathbf{e}_\dagger[\psi]$, $\mathbf{e}_\ddagger[\psi]$ denote assigned variables, unobserved variables and evidence in $\psi \in \Psi$. Then the query μ to P can be computed as follows:

$$\frac{\sum_{\psi \in \Psi} \left(\prod_{u_i \in \mathbf{x}[\psi] \cup \mathbf{z}_\dagger[\psi] \cup \mathbf{e}_\dagger[\psi]} P(u_i \mid \mathbf{ppa}(U_i)) f(\mathbf{x}[\psi]) R[\psi] \right)}{\sum_{\psi \in \Psi} \left(\prod_{u_i \in \mathbf{x}[\psi] \cup \mathbf{z}_\dagger[\psi] \cup \mathbf{e}_\dagger[\psi]} P(u_i \mid \mathbf{ppa}(U_i)) R[\psi] \right)} \quad (6)$$

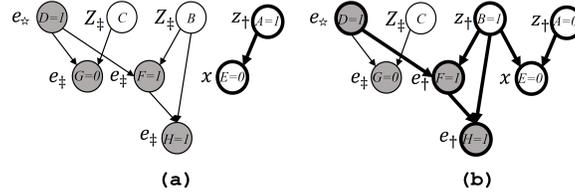


Figure 4: Two safe contextual assignments to RVs of BN of Figure 1a: (a) in the context $A = 1$, where edges $C \rightarrow E$ and $B \rightarrow E$ are redundant since $E \perp B, C \mid A = 1$; (b) in the context $A = 0, B = 1$, where the edge $C \rightarrow E$ is redundant since $E \perp C \mid A = 0, B = 1$. To identify such assignments, intuitively, we should apply the Bayes-ball algorithm after removing these edges. Portions of graphs that the algorithm visits, starting with visiting the variable E from its child, are highlighted. Notice that variables $\mathbf{X}, \mathbf{Z}_\dagger, \mathbf{E}_\dagger$ lie in the highlighted portion.

where $R[\psi]$ denotes $\mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[\psi]}]$.

We draw some important conclusions: i) μ can be exactly computed by performing the summation over safe contextual assignments; notably, variables in \mathbf{Z}_\dagger vary, and so do variables in \mathbf{E}_\dagger ; ii) For all $\psi \in \Psi$, the computation of $\mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[\psi]}]$ does not depend on the context $\mathbf{x}[\psi], \mathbf{z}_\dagger[\psi]$ since no basis of $\mathbf{e}_\dagger[\psi]$ is assigned in the context (by Theorem 5). Hence, $\mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[\psi]}]$ can be computed independently. However, the context decides which evidence should be in the subset $\mathbf{e}_\dagger[\psi]$. That is why we can not cancel $\mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[\psi]}]$ from the numerator and denominator.

5.2.2 CONTEXT-SPECIFIC LIKELIHOOD WEIGHTING

First, we present an algorithm that simulates a $\text{DC}(\mathcal{B})$ program \mathbb{P} and generates safe contextual assignments. Then we discuss how to estimate the expectations independently before estimating μ .

Simulation of $\text{DC}(\mathcal{B})$ Programs We start by asking a question. Suppose we modify the first and the fourth rule of Bayes-ball simulation, discussed in Section 3.2, as follows:

- In the first rule, when the visit of an unobserved variable is from its child, everything remains the same except that only some parents are visited, not all.
- Similarly, in the fourth rule, when the visit of an observed variable is from its parent, everything remains the same except that only some parents are visited.

Which variables will be assigned, and which will be weighted using the modified simulation rules? Intuitively, only a subset of variables in \mathbf{Z}_\star should be assigned, and only a subset of variables in \mathbf{E}_\star should be weighted. But then how to assign/weigh a variable knowing the state of only some of its parent. We can do that when structures are present within CPDs of \mathcal{B} , and these structures are explicitly represented using clauses in a $\text{DC}(\mathcal{B})$ program \mathbb{P} . Recall that the proof procedure, discussed in Section 5.1, can sample random variables without sampling some of their ancestors. Hence, the key idea is to visit only some parents

Algorithm 2 Simulation of $DC(\mathcal{B})$ Programs

procedure SIMULATE-GROUND-DC(\mathbf{x}, \mathbf{e})

- Simulates a $DC(\mathcal{B})$ program \mathbb{P} based on inputs: i) \mathbf{x} : a query; ii) \mathbf{e} : evidence.
 - Let $dep(\mathbb{P})$ be the dependency set of \mathbb{P} .
 - The procedure maintains global data structures: i) **Asg**, a table that records assignments of variables ($\mathbf{x} \cup \mathbf{z}_\dagger$); ii) **Forward**, a set of variables whose children to be visited from parent; iii) **Top**, a set of variables marked on top; iv) **Bottom**, a set of variables marked on bottom.
 - Output: i) $f(\mathbf{x})$ that can be either 0 or 1; ii) **W**: a table of weights of diagnostic evidence (\mathbf{e}_\dagger).
1. Empty **Asg**, **W**, **Top**, **Bottom**, **Forward**.
 2. If **PROVE-MARKED-GROUND**(\mathbf{x}) is **yes** then $f(\mathbf{x}) := 1$ else $f(\mathbf{x}) := 0$.
 3. While **Forward** is not empty:
 - (a) Remove **A** from **Forward** and add **A** to **Bottom**.
 - (b) For all A_0 such that $dep(\mathbb{P}) \models pa(A_0, A)$:
 - i. If A_0 is observed with value v in \mathbf{e} and $A_0 \notin \text{Top}$:
 - A. **Choose** $A_0 \sim D \leftarrow A_1, \dots, A_n \in \mathbb{P}$ such that **PROVE-MARKED-GROUND**(A_1, \dots, A_n) is **yes**. Add A_0 to **Top**.
 - B. Let p be the likelihood at v according to distribution D . Record $W[A_0] := p$.
 - ii. If A_0 is not observed in \mathbf{e} and $A_0 \notin \text{Bottom}$: add A_0 to **Forward**.
 4. Return $[f(\mathbf{x}), \mathbf{W}]$.
-

(if possible due to structures); consequently, those unobserved parents that are not visited might not be needed to be sampled.

To realize that, we need to adapt the Bayes-ball simulation such that it works on $DC(\mathcal{B})$ programs. However, there is a problem: \mathbb{P} is a set of clauses, and no explicit graph is associated with it on which the Bayes-ball can be applied. Fortunately, we can infer direct influence relationships using the dependency set of \mathbb{P} , which is automatically constructed, as discussed in Section 4.2. The adapted simulation for $DC(\mathcal{B})$ programs is defined procedurally in Algorithm 2. The algorithm visits variables from their parents and calls the top-down proof procedure (Algorithm 3) to visit variables from their children. Like Bayes-ball, these algorithms also mark variables on top and bottom to avoid repeating the same action.

Since the simulation of \mathbb{P} follows the same four rules of Bayes-ball simulation except that only some parents are visited in the first and fourth rule, we show that

Lemma 7. *Let \mathbf{E}_\dagger be a set of observed variables weighed and let \mathbf{Z}_\dagger be a set of unobserved variables, apart from query variables, assigned in a simulation of \mathbb{P} , then,*

$$\mathbf{Z}_\dagger \subseteq \mathbf{Z}_\star \text{ and } \mathbf{E}_\dagger \subseteq \mathbf{E}_\star.$$

Algorithm 3 DC(\mathcal{B}) Proof Procedure Marked

procedure PROVE-MARKED-GROUND(G)

- Proves a conjunction of ground atoms G , consequently, visits variables from child.
 - Accesses \mathbb{P} , **Top**, **Bottom**, **Forward**, **Asg**, and \mathbf{e} as defined in Algorithm 2.
 - Returns **yes** if there is a choice that makes G empty; otherwise the procedure fails.
1. While G is not empty:
 - (a) Select the first atom A from G .
 - (b) If A is of the form $\mathbf{t} \cong \mathbf{v}$:
 - i. If \mathbf{t} is observed in \mathbf{e} and its observed value is \mathbf{v} : remove A from G .
 - ii. Else if $\mathbf{t} \in \mathbf{Top}$:
 - A. If $\mathbf{Asg}[\mathbf{t}] = \mathbf{v}$: remove A from G .
 - iii. Else:
 - A. Add \mathbf{t} to **Top**.
 - B. **Choose** $\mathbf{t} \sim D \leftarrow B_1, \dots, B_n \in \mathbb{P}$ with \mathbf{t} in the head.
 - C. Set $G := B_1, \dots, B_n, \mathbf{sample}(\mathbf{t}, D), G$.
 - (c) Else If A is of the form $\mathbf{sample}(\mathbf{t}, D)$:
 - i. Sample a value \mathbf{v} from D , record $\mathbf{Asg}[\mathbf{t}] := \mathbf{v}$, and remove A from G .
 - ii. If $\mathbf{t} \notin \mathbf{Bottom}$: add \mathbf{t} to **Forward**
 2. Return **yes**.
-

Algorithm 4 Generation of residual evidence's weights for DC(\mathcal{B}) Programs

procedure WEIGHT-RES-GROUND(\mathbf{Res})

- Generates residual evidence's weights based on input: i) \mathbf{Res} , a list of residual evidence. This procedure accesses \mathbb{P} of Algorithm 2.
 - Output: i) W_1 , a table of residual evidence's weights.
1. For all random variables A_0 in \mathbf{Res} .
 - (a) Let \mathbf{v} be the value observed for A_0 in \mathbf{Res} .
 - i. **Choose** $A_0 \sim D \leftarrow A_1, \dots, A_n \in \mathbb{P}$ such that $\mathbf{PROVE-MARKED-GROUND}(A_1, \dots, A_n)$ is **yes**.
 - ii. Let \mathbf{p} be the likelihood of \mathbf{v} according to distribution D . Record $W_1[A_0] := \mathbf{p}$.
 2. Return W_1 .
-

The query variables \mathbf{X} are always assigned since the simulation starts with visiting these variables as if visits are from one of their children. To simplify notation, from now on we use \mathbf{Z}_\dagger to denote the subset of variables in \mathbf{Z}_\star that are assigned, \mathbf{E}_\dagger to denote the subset of variables in \mathbf{E}_\star that are weighted in the simulation of \mathbb{P} . \mathbf{Z}_\ddagger to denote $\mathbf{Z}_\star \setminus \mathbf{Z}_\dagger$, and \mathbf{E}_\ddagger

to denote $\mathbf{E}_\star \setminus \mathbf{E}_\dagger$. We show that the simulation performs safe contextual assignments to requisite variables.

Theorem 8. *Partial assignments \mathbf{x} , \mathbf{z}_\dagger , \mathbf{Z}_\dagger , \mathbf{e}_\dagger , \mathbf{e}_\ddagger generated in simulations of $DC(\mathcal{B})$ programs are safe contextual assignments.*

The proof of Theorem 8 relies on the following Lemma.

Lemma 9. *Let \mathbb{P} be a $DC(\mathcal{B})$ program specifying a distribution P . Let \mathbf{B}, \mathbf{C} be disjoint sets of parents of a variable A . In a simulation of \mathbb{P} , if A is sampled/weighted, given an assignment \mathbf{c} , and without assigning \mathbf{B} , then,*

$$P(A \mid \mathbf{c}, \mathbf{B}) = P(A \mid \mathbf{c}).$$

Furthermore, two contextual assignments generated in two simulations of a $DC(\mathcal{B})$ program are either identical or mutually exclusive. This is because the algorithm behind the two simulations is the same, so both will generate identical contextual assignments unless a different value is sampled for at least one unobserved variable. However, the two contextual assignments will clearly become mutually exclusive when a different value is sampled for an unobserved variable.

Hence, just like the standard LW, we sample from a factor Q_\dagger of the proposal distribution Q_\star , which is given by,

$$Q_\dagger = \prod_{u_i \in \mathbf{X} \cup \mathbf{Z}_\dagger \cup \mathbf{e}_\dagger} P(u_i \mid \mathbf{ppa}(U_i))$$

where $P(u_i \mid \mathbf{ppa}(U_i)) = 1$ if $u_i \in \mathbf{e}_\dagger$. It is precisely this factor that Algorithm 2 considers for the simulation of \mathbb{P} . Starting by first setting $\mathbf{E}_\star, \mathbf{E}_\dagger$ their observed values, it assigns $\mathbf{X} \cup \mathbf{Z}_\dagger$ and weighs \mathbf{e}_\dagger in the topological ordering. In this process, it records *partial weights* $\mathbf{w}_{\mathbf{e}_\dagger}$, such that: $\prod_{x_i \in \mathbf{e}_\dagger} w_{x_i} = w_{\mathbf{e}_\dagger}$ and $w_{x_i} \in \mathbf{w}_{\mathbf{e}_\dagger}$. Given M partially weighted samples $\mathcal{D}_\dagger = \langle \mathbf{x}[1], \mathbf{w}_{\mathbf{e}_\dagger[1]} \rangle, \dots, \langle \mathbf{x}[M], \mathbf{w}_{\mathbf{e}_\dagger[M]} \rangle$ from Q_\dagger , we could estimate μ using Theorem 6 as follows:

$$\bar{\mu} = \frac{\sum_{m=1}^M f(\mathbf{x}[m]) \times w_{\mathbf{e}_\dagger[m]} \times \mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[m]}]}{\sum_{m=1}^M w_{\mathbf{e}_\dagger[m]} \times \mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[m]}]} \quad (7)$$

However, we still can not estimate it since we still do not have expectations $\mathbb{E}_{Q_\star}[W_{\mathbf{e}_\dagger[m]}]$. Fortunately, there are ways to estimate them from partial weights in \mathcal{D}_\dagger . We discuss one such way next.

Estimating the Expected Weight of Residuals We start with the notion of sample mean. Let $\mathcal{W}_\star = \langle w_{e_1}[1], \dots, w_{e_m}[1] \rangle, \dots, \langle w_{e_1}[n], \dots, w_{e_m}[n] \rangle$ be a data set of n observations of weights of m diagnostic evidence drawn using the standard LW. How can we estimate the expectation $\mathbb{E}_{Q_\star}[W_{e_i}]$ from \mathcal{W}_\star ? The standard approach is to use the sample mean: $\overline{W}_{e_i} = \frac{1}{n} \sum_{r=1}^n w_{e_i}[r]$. In general, $\mathbb{E}_{Q_\star}[W_{e_i} \dots W_{e_j}]$ can be estimated using the estimator: $\overline{W}_{e_i} \dots \overline{W}_{e_j} = \frac{1}{n} \sum_{r=1}^n w_{e_i}[r] \dots w_{e_j}[r]$. Since LW draws are independent and identical distributed (i.i.d.), it is easy to show that the estimator is unbiased.

However, some entries, i.e., weights of residual evidence, are missing in the data set \mathcal{W}_\dagger obtained using CS-LW. The trick is to fill the missing entries by drawing samples of the

missing weights once we obtain \mathcal{W}_\dagger . More precisely, missing weights $\langle W_{e_i}, \dots, W_{e_j} \rangle$ in r^{th} row of \mathcal{W}_\dagger are filled in with a joint state $\langle w_{e_i}[r], \dots, w_{e_j}[r] \rangle$ of the weights. To draw the joint state, we use Algorithm 4 and call `WEIGHT-RES-GROUND`($[e_i, \dots, e_j]$) to visit observed variables $\langle E_i, \dots, E_j \rangle$ from parent. Once all missing entries are filled in, we can estimate $\mathbb{E}_{Q_\star}[W_{e_i} \dots W_{e_j}]$ using the estimator $\overline{W_{e_i} \dots W_{e_j}}$ as just discussed. Once we estimate all required expectations, it is straightforward to estimate μ using Equation 7.

Interpretation At this point, we can gain some insight into the role of CSIs in sampling. They allow us to estimate the expectation $\mathbb{E}_{Q_\star}[W_{e_\dagger}]$ separately. We estimate it from all samples obtained at the end of the sampling process, thereby reducing the contribution W_{e_\dagger} makes to the variance of our main estimator $\bar{\mu}$. The residual evidence e_\dagger would be large if many CSIs are present in the distribution; consequently, we would obtain a much better estimate of μ using significantly fewer samples. Moreover, drawing a single sample would be faster since only a subset of requisite variables is visited. Hence, in addition to CIs, we exploit CSIs and improve LW further. We observe all these speculated improvements in our experiments.

6. Inference in First-Order DC# Programs

This section extends CS-LW to first-order DC# programs where clauses need not be mutually exclusive. The extended algorithm is called *first-order context-specific likelihood weighting* (FO-CS-LW).

Using tools of logic such as unification and substitution, it is easy to simulate first-order programs without completely grounding them first. The simulation process defined in Algorithm 5 does that. It uses dependency sets of programs to visit RVs from parents and calls Algorithm 7 to visit RVs from children. The RV sets of programs are used to identify RVs defined by programs. There are two important features of the top-down proof procedure with logical variables defined in Algorithm 7, which are worth highlighting.

Firstly, it computes answer substitutions of a query $\leftarrow A_1, \dots, A_m$, that is, the substitutions of the refutations of the query restricted to the variables in the query. To realize that, instead of the query, the initial goal G_0 is of the form:

$$\text{yes}(V_1, \dots, V_k) \leftarrow A_1, \dots, A_m$$

where V_1, \dots, V_k are the logical variables that appear in the query. This allows the procedure to return the answer $\{V_1/\tau_1, \dots, V_k/\tau_k\}$ when the body of goal G_i is empty after applying the resolution rules. This is just like the proof procedure for definite programs with logical variables (Poole & Mackworth, 2010).

Secondly and more importantly, the procedure searches and collects all distributions defined for an RV in a global variable `Dst` before sampling values of the RV from the combined distribution obtained using the combining rules. This is as per the semantics of DC# programs, and in this way, the algorithm also exploits ICIs.

The generation of residual evidence’s weights is defined in Algorithm 7.

Dealing with Continuous RVs Till now, we have not talked about inference in programs with continuous RVs. Nevertheless, the notions discussed so far are sufficient to explain inference in such programs. Likelihood weighting naturally extends to continuous

Algorithm 5 Simulation of DC# Programs

procedure SIMULATE-DC#(\mathbf{x}, \mathbf{e})

- Simulates a DC program \mathbb{P} based on inputs: i) \mathbf{x} , a ground query; ii) \mathbf{e} : evidence.
 - Let $dep(\mathbb{P})$ be the dependency set of \mathbb{P} .
 - The procedure maintains global data structures: i) **Asg**, a table that records partial assignments of unobserved variables; ii) **Forward**, a set of variables whose children to be visited from parent; iii) **Top**, a set of variables marked on top; iv) **Bottom**, a set of variables marked on bottom; v) **Dst**, a table that records variables' distributions.
 - Output: i) $f(\mathbf{x})$ that can be either 0 or 1; ii) **W**: a table of weights of diagnostic evidence (\mathbf{e}_\dagger).
1. Empty **Asg**, **W**, **Top**, **Bottom**, **Dst**, **Forward**. Let \mathcal{CR} be combining rules used for \mathbb{P} .
 2. If **PROVE-MARKED**(\mathbf{x}) fails then $f(\mathbf{x}) := 0$ else $f(\mathbf{x}) := 1$
 3. While **Forward** is not empty:
 - (a) Remove **A** from **Forward** and add **A** to **Bottom**
 - (b) For all A_0 such that $dep(\mathbb{P}) \models pa(A_0, A)$:
 - i. If A_0 is observed to be v in \mathbf{e} and $A_0 \notin \text{Top}$:
 - A. Add A_0 to **Top**.
 - B. For all (renamed) clause $B_0 \sim D \leftarrow B_1, \dots, B_m \in \mathbb{P}$ such that mgu σ unifies B_0 and A_0 : call **PROVE-MARKED**($(B_1, \dots, B_m, \text{add_dst}(B_0, D))\sigma$)
 - C. Let $[D_1, \dots, D_m]$ be distributions for A_0 recorded in table **Dst**, and let p be the likelihood at v according to $\mathcal{CR}([D_1, \dots, D_m])$. Record $W[A_0] := p$.
 - ii. If A_0 is not observed in \mathbf{e} and $A_0 \notin \text{Bottom}$: add A_0 to **Forward**
 4. Return $[f(\mathbf{x}), \mathbf{W}]$
-

domains (Koller & Friedman, 2009). We just need rules to resolve comparison atoms of the form $V_1 \diamond V_2$ that appear in the body of clauses. These rules are already present in the proof procedure.

To ensure that FO-CS-LW correctly estimates the probabilities, we need to ensure that the simulation of DC# programs also generates safe contextual assignments that are identical or mutually exclusive. For this, we just need to ensure that Lemma 9 also holds for the DC# programs' simulation.

Theorem 10. *Lemma 9 is also true for simulation of DC# programs.*

7. Empirical Evaluation

In this section, we empirically evaluate our inference algorithms and answer several research questions.

Algorithm 6 DC# Proof Procedure

procedure PROVE-MARKED(Q)

- Proves a conjunction of atoms Q with variables V_1, \dots, V_k .
 - Accesses \mathbb{P} , **Asg**, **Top**, **Bottom**, **Dst**, **Forward**, **e**, and \mathcal{CR} as defined in Algorithm 5.
 - Returns substitutions of variables V_1, \dots, V_k if refutation exists; otherwise fails.
 - Let $rv(\mathbb{P})$ be an RV set of \mathbb{P} .
1. Set G to a clause $\text{yes}(V_1, \dots, V_k) \leftarrow Q$
 2. While the body of G is not empty:
 - (a) Suppose G is $\text{yes}(t_1, \dots, t_k) \leftarrow A_1, \dots, A_n$. Select A_1
 - (b) If A_1 is an atom of the form $X \cong V$:
 - i. **Choose** a grounding substitution σ_1 such that $rv(\mathbb{P}) \models X\sigma_1$
 - ii. If $X\sigma_1$ is observed in **e**: let v be the observed value.
 - iii. Else if $X\sigma_1 \in \text{Top}$: let v be $\text{Asg}[X\sigma_1]$.
 - iv. Else:
 - A. For all (renamed) clause $B_0 \sim D \leftarrow B_1, \dots, B_m \in \mathbb{P}$ such that mgu σ_2 unifies B_0 and $X\sigma_1$: call $\text{PROVE-MARKED}((B_1, \dots, B_m, \text{add_dst}(B_0, D))\sigma_2)$
 - B. Let $M = [D_1, \dots, D_1]$ be distributions for $X\sigma_1$ recorded in table **Dst**. Let v be value sampled from $\mathcal{CR}(M)$. Record $\text{Asg}[X\sigma_1] := v$.
 - C. Add $X\sigma_1$ to **Top**. If $X\sigma_1 \notin \text{Bottom}$ then add $X\sigma_1$ to **Forward**
 - v. If mgu σ_3 unify $X\sigma_1 \cong v$ and $X \cong V$: set $G := ((\text{yes}(t_1, \dots, t_k) \leftarrow A_2, \dots, A_n)\sigma_1)\sigma_3$
 - (c) Else if A_1 is of the form $\text{add_dst}(X, D)$: // A_1 will always be ground here
 - i. Record $\text{Dst}[X] := D$ and set $G := \text{yes}(t_1, \dots, t_k) \leftarrow A_2, \dots, A_n$
 - (d) Else if A_1 is of the form $V_1 \diamond V_2$ and evaluates to true: $G := \text{yes}(t_1, \dots, t_k) \leftarrow A_2, \dots, A_n$
 3. Return $\{V_1/t_1, \dots, V_k/t_k\}$ when G is $\text{yes}(t_1, \dots, t_k) \leftarrow$
-

7.1 How do the sampling speed and the accuracy of estimates obtained using CS-LW compare with the standard LW in the presence of CSIs?

To answer it, we need BNs with structures present within CPDs. Such BNs, however, are not readily available since the structure while designing inference algorithms is generally overlooked. We identified two BNs from the Bayesian network repository (Elidan, 2001), which have many structures within CPDs: i) *Alarm*, a monitoring system for patients with 37 variables; ii) *Andes*, an intelligent tutoring system with 223 variables.

We used the standard decision tree learning algorithm to detect structures and overfitted it on tabular-CPDs to get tree-CPDs, which was then converted into clauses. Let us denote the program with these clauses by \mathbb{P}_{tree} . CS-LW is implemented in the Prolog programming language, thus to compare the sampling speed of LW with CS-LW, we need a similar implementation of LW. Fortunately, we can use the same implementation of CS-LW for obtaining LW estimates. Recall that if we do not make structures explicit in clauses and represent each entry in tabular-CPDs with clauses, then CS-LW boils down to LW. Let \mathbb{P}_{table} denotes the program where each rule in it corresponds to an entry in tabular-CPDs. Table 1 shows the comparison of estimates obtained using \mathbb{P}_{tree} (CS-LW) and \mathbb{P}_{table} (LW). Note that CS-LW automatically discards non-requisite variables for sampling. So, we chose the

Algorithm 7 Generation of residual evidence’s weights for DC# Programs**procedure** WEIGHT-RESIDUALS(Res)

- Generates residual evidence’s weights based on input: i) Res, a list of residual evidence. This procedure accesses \mathbb{P} and \mathcal{CR} of Algorithm 5.
 - Output: i) W_1 , a table of residual evidence’s weights.
1. For all random variable A_0 in Res.
 - (a) Let v be the value observed for A_0 in Res.
 - i. For all (renamed) clause $B_0 \sim D \leftarrow B_1, \dots, B_m \in \mathbb{P}$ such that mgu σ unifies B_0 and A_0 : call PROVE-MARKED($(B_1, \dots, B_m, \text{add_dst}(B_0, D))\sigma$)
 - ii. Let $[D_1, \dots, D_1]$ be distributions for A_0 recorded in table Dst, and let p be the likelihood at v according to $\mathcal{CR}([D_1, \dots, D_1])$. Record $W_1[A_0] := p$.
 2. Return W_1 .

		LW		CS-LW	
BN	N	MAE \pm Std.	Time	MAE \pm Std.	Time
<i>Alarm</i>	100	0.2105 \pm 0.1372	0.09	0.0721 \pm 0.0983	0.06
	1000	0.0766 \pm 0.0608	0.86	0.0240 \pm 0.0182	0.53
	10000	0.0282 \pm 0.0181	8.64	0.0091 \pm 0.0069	5.53
	100000	0.0086 \pm 0.0067	89.93	0.0034 \pm 0.0027	57.64
<i>Andes</i>	100	0.0821 \pm 0.0477	1.07	0.0619 \pm 0.0453	0.22
	1000	0.0257 \pm 0.0184	10.62	0.0163 \pm 0.0139	2.20
	10000	0.0087 \pm 0.0069	106.55	0.0058 \pm 0.0042	22.62
	100000	0.0025 \pm 0.0015	1074.93	0.0020 \pm 0.0016	233.72

Table 1: The mean absolute error (MAE), the standard deviation of the error (Std.), and the average execution time (in seconds) versus the number of samples (N). For each case, LW and CS-LW were executed 30 times.

query and evidence such that almost all variables in BNs were requisite for the conditional query.

As expected, we observe that less time is required by CS-LW to generate the same number of samples. This is because it visits only the subset of requisite variables in each simulation. *Andes* has more structures compared to *Alarm*. Thus, the sampling speed of CS-LW is much faster compared to LW in *Andes*. Additionally, we observe that the estimate, with the same number of samples, obtained by CS-LW is much better than LW. This is significant. It is worth mentioning that approaches based on collapsed sampling obtain better estimates than LW with the same number of samples, but then the speed of drawing samples significantly decreases (Koller & Friedman, 2009). In CS-LW, the speed increases when structures are present. This is possible because CS-LW exploits CSIs.

Hence, we get the answer to our first question: When many structures are present, and when they are made explicit in clauses, then CS-LW will draw samples faster compared to LW. Additionally, estimates will be better with the same number of samples.

7.2 How does FO-CS-LW perform as the domain size increases?

The exact inference algorithms that PLP systems generally use for inference do not scale with domain sizes of logical variables in the program. Large domain sizes result in a huge ground program on which exact inference becomes intractable even for PLPs supporting only Boolean RVs. Thus, it is interesting to investigate how FO-CS-LW, an approximate inference algorithm, performs on such PLPs.

For this purpose, we compared FO-CS-LW with the inference algorithm used in ProbLog, one of the most popular PLP systems. This algorithm first grounds first-order programs and then performs exact inference to exploit structures of clauses in the programs. We used a ProbLog program shown in Figure 5 for this experiment. Note that in ProbLog, when multiple distributions are specified for an RV, they are combined using NoisyOR, and when no distribution is specified, the RV is set to false. So, using this combining rule, the ProbLog program of Figure 5 can be expressed as a DC# program. The domain size of each logical variable in the program is two since there are two clients, two accounts, and two loans. In such cases, instead of specifying the domain size of each logical variable separately, we will simply say that the domain size is two. Notice that relationships among clients, accounts, and loans are also probabilistic, so the number of RVs explicated by the program becomes huge as the domain size increases. More precisely, there are $3n^2 + 6n$ RVs when domain size is n .

We also compared the performance of FO-CS-LW when applied directly to equivalent first-order programs versus when applied to the grounded programs. Indeed, when the full ground network is huge, and the requisite network is also huge, it is better first to ground the programs and then apply FO-CS-LW. This is because unifications used to reason in first-order programs are somewhat costly operations, which are performed once if programs are grounded first. This case is illustrated in Figure 6, where the probability of query $Q_1 = P(\text{high_savings}(a1) \cong t \mid \text{home_loan}(l1) \cong f, \text{debt}(c1) \cong t, \text{has_loan}(c1, l1) \cong f)$ is estimated and FO-CS-LW performs well if programs are grounded first. However, if the full ground network is huge, but the requisite network is very small, it is better to reason on the first-order level. This is because the cost of searching a few relevant clauses in a huge set of ground clauses exceeds the unification cost. This is the case when the query (Q_2) is to compute $P(\text{debt}(c1) \cong t)$ given all other RVs (all RVs of type `has_loan(C, L)` were set to `f` (“false”), all RVs of type `high_savings(A)` were set to `f`, and rest RVs were set to `t` (“true”). Furthermore, ProbLog could not compute probabilities beyond the domain size of 9 on a machine with 132 GB main memory, whereas FO-CS-LW quickly scaled to a domain size of 50.

One might expect that FO-CS-LW would perform poorly as domain size increases, and more samples would be required to get a reasonable estimate of probabilities. Figure 6 suggests the opposite. Using the same number of samples, the standard deviation from the mean does not increase as domain size increases. This is because FO-CS-LW exploits symmetries that arise due to noisy OR. Notice that exact probabilities do not change much

<pre> client(c1). client(c2). ... account(a1). account(a2). ... loan(l1). loan(l2). ... 0.7 :: home_loan(L):- loan(L). 0.3 :: high_savings(A):- account(A). 0.01 :: has_account(C, A):- client(C), account(A). 0.02 :: account_loan(A, L):- account(A), loan(L). 0.9 :: has_loan(C, L):- has_account(C, A), account_loan(A, L). 0.001 :: has_loan(C, L):- client(C), loan(L). 0.9 :: debt(C):- has_loan(C, L), home_loan(L). 0.6 :: debt(C):- has_loan(C, L), \+home_loan(L). 0.3 :: debt(C):- has_account(C, A), \+high_savings(A). 0.01 :: debt(C):- client(C). </pre>	<pre> client(c1) ~ val(t). client(c2) ~ val(t). ... account(a1) ~ val(t). account(a2) ~ val(t). ... loan(l1) ~ val(t). loan(l2) ~ val(t). ... home_loan(L) ~ bernoulli(0.7) ← loan(L) ≅ t. high_savings(A) ~ bernoulli(0.3) ← account(A) ≅ t. has_account(C, A) ~ bernoulli(0.01) ← client(C) ≅ t, account(A) ≅ t. account_loan(A, L) ~ bernoulli(0.02) ← account(A) ≅ t, loan(L) ≅ t. has_loan(C, L) ~ bernoulli(0.9) ← has_account(C, A) ≅ t, account_loan(A, L) ≅ t. has_loan(C, L) ~ bernoulli(0.001) ← client(C) ≅ t, loan(L) ≅ t. debt(C) ~ bernoulli(0.9) ← has_loan(C, L) ≅ t, home_loan(L) ≅ t. debt(C) ~ bernoulli(0.6) ← has_loan(C, L) ≅ t, home_loan(L) ≅ f. debt(C) ~ bernoulli(0.3) ← has_account(C, A) ≅ t, high_savings(A) ≅ f. debt(C) ~ bernoulli(0.01) ← client(C) ≅ t. </pre>
--	--

Figure 5: (Left) A ProbLog program specifying a probability distribution over relationships and attributes of clients, accounts, and loans. (Right) The equivalent DC# program.

as domain size increases because unique parameters do not increase even though the domain size increases and the number of parameters increases.

We conclude that FO-CS-LW scales with the domain size and can be useful on problems where the ProbLog inference algorithm fails.

7.3 How does FO-CS-LW compare to the state-of-the-art inference algorithms for hybrid relational probabilistic models?

Those exact inference algorithms that exploit CSIs are not readily applicable to hybrid relational probabilistic models. There are not many open-source implementations of inference algorithms for such models⁵. The exploitation of first-order CSIs and symmetries arising due to aggregations in such models is challenging, and to the best of our knowledge, no algorithm can exploit them in such models. To some extent, the likelihood weighting-based inference algorithm developed for the old version of DC can exploit them (Nitti et al., 2016). However, it does not filter out irrelevant evidence before inference, which is crucial in the case of relational databases. So, in this experiment, we aim to investigate how much FO-CS-LW improves upon old-DC’s inference algorithm.

For this purpose, we used a real-world relational data generated by processing the financial database from the PKDD’99 Discovery Challenge. This data set is about services that a bank offers to its clients, such as loans, accounts, and credit cards. It contains information of four types of entities: 5,358 clients, 4,490 accounts, 680 loans and 77 districts. Ten attributes are of the continuous type, and three are of the discrete type. The data set contains four relations: `hasAccount/2` that links clients to accounts; `hasLoan/2` that links accounts to loans; `clientDistrict/2` that links clients to districts; and finally `clientLoan/2` that links clients to loans.

5. The publicly available code for BLPs is based on SICStus Prolog 3, the version that SICStus no longer supports. So, it is difficult to run the code.

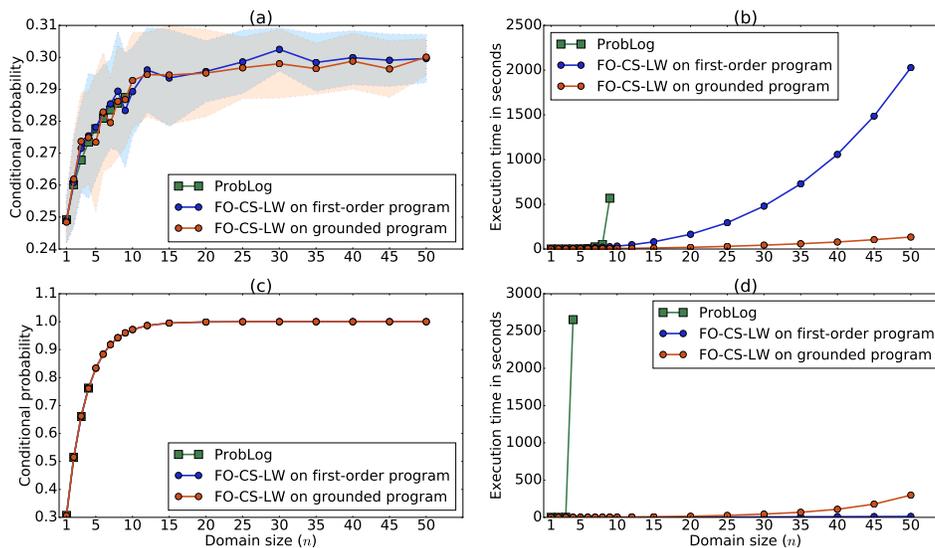


Figure 6: The first two graphs are for query Q_1 , and the last two graphs are for query Q_2 , mentioned in Section 7.2. (a) Comparison of exact probabilities of Q_1 computed by ProbLog, and probabilities estimated by FO-CS-LW when applied to the first-order and the grounded program. (b) The average time for processing Q_1 versus the domain size. (c) Comparison of probabilities of query Q_2 for the three cases. (d) The average time for processing Q_2 versus the domain size. To estimate the probabilities, FO-CS-LW used 10,000 samples. The size of n indicates that n clients, n accounts, and n loans are present in the program. The shaded region denotes the standard deviation from the mean probability estimated by FO-CS-LW when executed 30 times for each n . The execution time includes the grounding time when FO-CS-LW is applied to grounded programs.

We learned a model in the form of distributional clauses as described in (Kumar, Kuželka, & De Raedt, 2021) from the financial data set. The learned model, which was a program, specified a probability distribution over all attributes of all instances of the entities in the data set. Since the program was learned just as described in (Kumar et al., 2021), relationships among entities could not be probabilistic. Furthermore, clauses in the program were mutually exclusive, and *aggregation atoms* and *statistical model atoms* were used in the bodies of the clauses. *Negations* were allowed in the bodies to deal with missing values or missing relationships. Details about these advanced constructs are present in Appendix B. A snippet of the learned program is shown in Figure 9.

Next, we created multiple subsets of the financial data set with varying numbers of accounts. These subsets were created by considering `account` to be the central entity. All information about clients, loans, and districts related to an account appeared in the same subset. All subsets had an account `a_10001` that was linked with a loan `l_7034`. Two queries to

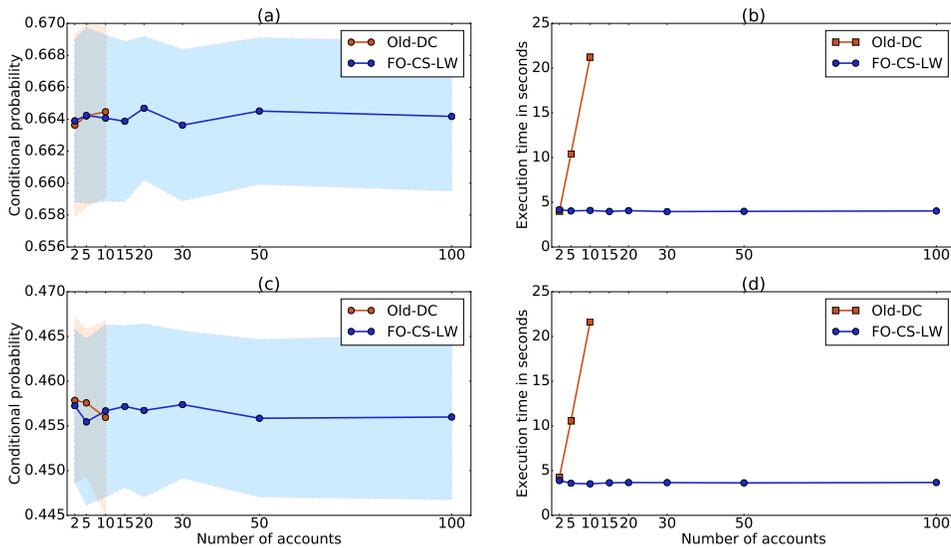


Figure 7: The first two graphs are for query Q_3 , and the last two graphs are for query Q_4 , mentioned in Section 7.3. (a) Comparison of estimates obtained using FO-CS-LW and Old-DC for query Q_3 . (b) Average time taken by FO-CS-LW and Old-DC for Q_3 . (c) Comparison of estimates for query Q_4 . (d) Average time taken for Q_4 . 10,000 samples were used for each of the two algorithms. The shaded region denotes the standard deviation from the mean probability estimated by algorithms when executed 100 times for each case. Old-DC suffered arithmetic underflow as observed data increased.

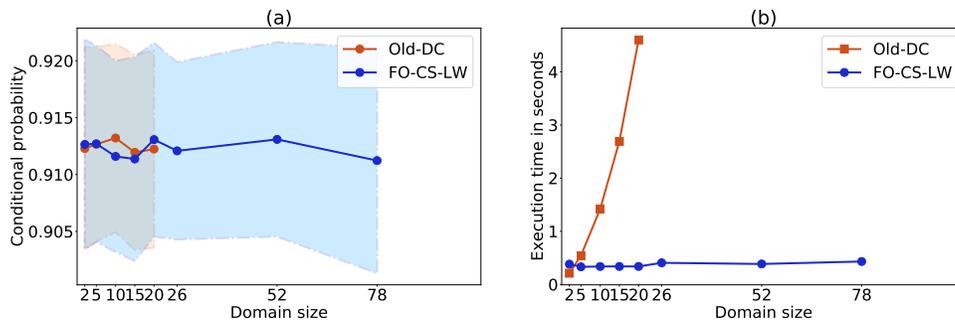


Figure 8: (a) Comparison of estimates obtained using FO-CS-LW and Old-DC for query Q_5 . (b) Average time taken by FO-CS-LW and Old-DC for Q_5 .

the learned program were considered: i) $Q_3 = P(\text{stdMonthInc}(a_{10001}) \cong X, X > 22000.0)$ given the rest subsets of data, ii) $Q_4 = P(\text{loanAmount}(1.7034) \cong X, X > 20000.0)$ given the

```

%Entities
account(a_10001) ~ val(true).
loan(l_7034) ~ val(true).
client(c_12303) ~ val(true).
district(d_24) ~ val(true).
...
%Relations
clientLoan(c_12303, l_7034) ~ val(true).
hasLoan(a_10001, l_7034) ~ val(true).
hasAccount(c_12303, a_10001) ~ val(true).
clientDistrict(c_12303, d_24) ~ val(true).
...
%Model
loanAmount(L) ~ gaussian(151080.0, 12853887266.5) ← loan(L) ≅ true.
...
stdMonthInc(M) ~ gaussian(M, 36225870.62) ← account(A) ≅ true, avgSumOfW(A) ≅ X, avgNrWith(A) ≅ Y, avg(Z1, (hasLoan(A, L) ≅ true,
loanAmount(L) ≅ Z1), Z), freq(A) ≅ m, linear([X, Y, Z], [0.34, -516.12, 0.004, 12797.18], M).

```

Figure 9: A snippet of the program learned from the financial data set. The aggregation atom `avg(-, -, -)` in the body of the last clause collects all loans’ amounts that a client has into a list and unifies `Z` with the average of the list. The statistical model atom `linear(-, -, -)` implements a linear function `M` is $0.34 \times X - 516.12 \times Y + 0.004 \times Z + 12797.18$. Predicates like `loanAmount/1`, `stdMonthInc/1`, `...`, represent the attributes of entities.

rest subsets of data. Figure 7 shows the comparison of estimates obtained for queries Q_3 and Q_4 . The old-DC suffered arithmetic underflow as observed data increased, whereas FO-CS-LW scaled quite well. An important observation, in this case, is that the execution time of FO-CS-LW does not increase by adding more data. This is because real-world data is often not highly relational, e.g., clients generally have one or two accounts, not ten accounts. This is the case here. By adding more data, we are just adding data irrelevant to account `a_10`. FO-CS-LW that exploits CIs can detect that, but the old-DC inference engine can not.

Similar observations were made when we used the program learned from the real-world NBA data set as described in (Kumar et al., 2021). This data set is about basketball matches from the National Basketball Association (Schulte & Routley, 2014). It records information about matches played between two teams and actions performed by each player of those two teams in the matches. This data set also contains relations and discrete-continuous attributes. We created multiple subsets of the data set with varying numbers of actions (the domain size). These subsets were created by considering actions to be the central entity. All subsets had actions of a player with id 41 in game 27. The query that we considered was: $Q_5 = P(\text{points}(27, 41) \cong X, X > 10)$ given the rest subsets of data. As shown in Figure 8, the observations are similar to those made for the financial dataset.

Thus, we conclude that when a significant amount of independencies are present in hybrid relational probabilistic models, FO-CS-LW outperforms the state-of-the-art inference algorithm of Old-DC.

8. Related Work

We describe relationships of the $DC\#$ framework introduced here with previous works in terms of representation and inference.

8.1 Relation to Representation Languages

Several relational representation languages based on aggregation functions and combining rules have been introduced in the past. Examples include probabilistic relational models (Friedman, et al., 1999, PRMs), directed acyclic probabilistic entity-relationship models (Heckerman, Meek, & Koller, 2004, DAPER), probabilistic relational language (Getoor & Grant, 2006, PRL), Bayesian logic programs (Kersting & De Raedt, 2007, BLPs), and first-order conditional influence language (Natarajan et al., 2008, FOCIL). These languages consist of two components: a qualitative component representing the relational structure of the domain (either using graphs or definite clauses) and a quantitative component specifying conditional probability distributions (CPDs) in the tabular form. So, they do not qualitatively represent the structures present within the CPDs.

A tree or a rule-based representation lets us represent the structures qualitatively (Boutilier et al., 1996; Poole, 1997; Ngo & Haddawy, 1997). That is why, probabilistic logic programming (PLP) has received much attention. Over the past three decades, many PLP languages have been proposed: PHA (Poole, 1993), Prism (Sato & Kameya, 1997), LPADs (Vennekens, Verbaeten, & Bruynooghe, 2004), ProbLog (De Raedt et al., 2007), ICL (Poole, 2008), CP-Logic (Vennekens, Denecker, & Bruynooghe, 2009). However, only a few of them support both discrete and continuous RVs: HProbLog (Gutmann et al., 2010), DC (Gutmann et al., 2011; Nitti et al., 2016), Extended-Prism (Islam et al., 2012), Hybrid-cplint (Alberti et al., 2017), (Michels et al., 2016). Nonetheless, the hybrid PLPs that have been studied in the past do not support combining rules, which is a core component of PLPs. BLPs do support combining rules, but then they do not qualitatively represent the structures. The syntax of $DC\#$ is the same as the syntax of DC introduced by (Nitti et al., 2016), but the semantics is different as $DC\#$ supports combining rules. The semantics of $DC\#$ is based on BLPs, so one can view $DC\#$ programs as BLPs qualitatively representing structures. Furthermore, a huge fragment of Problog programs, one of the most popular PLPs, is expressible as $DC\#$ programs. Annotated disjunctions and directed cycles are not supported in $DC\#$ currently.

$DC\#$ is a probabilistic programming language (PPL); thus, we should also describe how $DC\#$ relates to various other PPLs that generally extends imperative programming with probability distributions such as Infer.NET (Minka, et al., 2010), BLOG (Li & Russell, 2013), Stan (Team, 2015), etc. $DC\#$ relates to these languages in the manner similar to how logic programming relates to imperative programming. One must use *if-then conditions* in these languages to express the structures and *for loops* to express aggregations/combining rules. Using statistical model atoms in the body of DCs, one can describe complex probabilistic models written in these languages. However, $DC\#$ currently does not support writing open-universe models that some PPLs, such as BLOG, support. On the other hand, it is unclear how to write DC programs with negations in these PPLs, which is important when dealing with missing data.

8.2 Relation to Inference Algorithms

Resolving representation issues does not quickly imply that crucial issues inherent to inference are also resolved. The prime focus of past works on relational representation languages, for example, those discussed at the beginning of the previous section, has been the development of syntax and semantics so that these languages succinctly represent first-order probabilistic models. For inference, most of them construct ground BNs and rely on inference algorithms for BNs. Thus, they do not exploit symmetrical parameters that arise due to grounding the first-order models. Indeed they do not exploit structures present within CPDs of ground BNs since they do not even qualitatively represent them.

The case of the related PPLs is quite surprising. A remarkable feature of PPLs is that the structures can be represented using if-then conditions in programs. However, equivalent BNs are constructed during inference, and algorithms for BNs are used; thus, CSIs implied by the structures are ignored. For example, Stan, the most popular PPLs that is also commercially used, does that. Interestingly, BLOG (Milch, et al., 2005) does exploit CSIs, but a different class of CSIs, CSIs that arise in open-universe models due to the addition and subtraction of some RVs from the model when some other RVs take certain values. It does not exploit those that are implied by if-then conditions even in closed-universe models/programs. The exploitation of such CSIs and symmetrical parameters that appear after unrolling (similar to grounding in relational models) for loops even in closed-universe PPLs are not well studied.

Leaving expressive languages aside, the exploitation of CSIs implied by structures within CPDs of ground BNs is itself a difficult problem that has puzzled researchers for decades. Research in this direction has mainly been focused on exact inference (Boutilier et al., 1996; Poole & Zhang, 2003). Nowadays, it is common to use knowledge compilation-based exact inference for this purpose (Chavira & Darwiche, 2008; Fierens et al., 2015; Shen, Choi, & Darwiche, 2016); however, this approach does not apply to hybrid models. The problem of exploiting CSIs in hybrid models is non-trivial and is poorly studied. Recently, it has attracted some attention (Zeng & Van Den Broeck, 2020). However, proposed approaches are also exact and rely on complicated weighted model integration (Belle, Passerini, & Van Den Broeck, 2015), which do not scale (Feldstein & Belle, 2021). CS-LW that we propose for BNs is simple, scalable, and applies to hybrid BNs.

FO-CS-LW that extends CS-LW to first-order DC# programs also exploits symmetries in such programs. So, it makes sense to relate FO-CS-LW with the literature of lifted inference algorithms that aim to exploit symmetries in first-order models. Broadly these algorithms can be divided into two categories. Firstly, those that exploit symmetries but do not exploit CSIs. Belonging to this category, there are exact (Poole, 2003; De Salvo Braz, Amir, & Roth, 2006; Getoor & Taskar, 2007; Kisynski & Poole, 2009; Taghipour, et al., 2013) and approximate algorithms (Niepert, 2012; Ahmadi, et al., 2013; Chen, et al., 2020). Belonging to the second category are those that also exploit CSIs. These algorithms are generally exact and do not easily apply to hybrid models (Van Den Broeck et al., 2011; Gogate & Domingos, 2011). Recently, an exact inference algorithm belonging to the second category was applied to hybrid models (Feldstein & Belle, 2021). However, this algorithm inherits the limitations of exact lifted inference. It is well known that exact lifted inference has two major limitations: i) only a small class of models is liftable (Van Den Broeck, 2011;

Kazemi, et al., 2016); ii) the inference is hard when binary relations are observed (Van Den Broeck & Darwiche, 2013). An any-time lifted inference algorithm was developed by (de Salvo Braz, et al., 2009), which provides intervals (upper and lower bounds) in which the desired answer is guaranteed to lie. However, even this is an exact method, and it may not scale well when one requires the intervals to be narrow. FO-CS-LW belongs to the second category of algorithms, which is approximate thus does not have these limitations. It applies to DC# language in which a larger class of models can be written, which is remarkable.

9. Future Work

Our paper intends to introduce an interesting probabilistic programming framework that can be further developed in the future. Several features can be added to this framework, which are discussed next.

The semantics of DC# can be extended to support open-universe models. The syntax already supports them, which was shown at the beginning of Section 4.2. One can borrow ideas from BLOG for this purpose (Milch & Russell, 2009). However, extending FO-CS-LW to the language supporting open-universe models will require solving a complex problem: how to detect requisite networks in such models while sampling? To detect them in closed-universe models, RV and dependency sets are automatically constructed by static analysis of programs as discussed in this paper; however, it is difficult to construct such sets for open-universe models as it might require dynamic analysis.

Our algorithm does not support inference in DC# programs specifying mixtures of discrete and continuous distributions. This is the case when some clauses state that an RV is distributed according to continuous distributions, while others state that the same RV is distributed according to discrete distributions. More precisely, suppose a program contains the following two clauses:

```
credit_score(C) ~ gaussian(700, 10.9) ← has_loan(C, L) ≅ true.
credit_score(C) ~ discrete([0.9 : 600, 0.1 : 650]) ← has_loan(C, L) ≅ false.
```

It is not clear whether the client’s credit score is of discrete type or continuous type. Likelihood weighting (LW) based algorithms do not estimate correct probabilities in such a case. To resolve this issue, (Wu, et al., 2018) proposed lexicographic LW. A context-specific variant of the lexicographic LW is needed for such programs. This has been left for the future.

As shown in Example 7, the `val(.)` functor in DC# provides an elegant way of describing deterministic dependencies (constraints) among RVs in probabilistic models (Mateescu & Dechter, 2009). It is, however, well known that sampling algorithms, in general, perform poorly when a significant amount of determinism is present in models. So, an important and interesting avenue for future work is to combine constraint propagation with FO-CS-LW.

Finally, we aim to open up a new direction towards improved sampling algorithms that exploit CSIs. Like LW, we believe MCMC algorithms can also be extended along the same line.

10. Conclusion

We have introduced a hybrid PLP framework called DC#. It supports combining rules required to describe relational models succinctly.

We have emphasized the exploitation of both CIs and CSIs for efficient inference. After realizing that a sampling algorithm that can properly do that is not well studied, we studied the role of CSIs in sampling. Subsequently, we introduced a notion of contextual assignment to show that CSIs allow for breaking the main problem of estimating conditional probability queries into several small problems that can be estimated independently. Based on this notion, we presented CS-LW that exploits CSIs implied by structures present within CPDs of BNs. We empirically showed that when a significant amount of structures are present, CS-LW generates samples faster than the standard LW and provides a better estimate of the query with much fewer samples.

Since CS-LW is based on theorem proving, we easily extended it with unification and substitution of logical variables for inference in first-order DC# programs. The resulting first-order inference algorithm, called FO-CS-LW, naturally exploits symmetries or ICIs that arise due to combining rules in programs.

Empirical results showed that FO-CS-LW scales with the domain size and outperforms the inference algorithm of the state-of-the-art hybrid PLP.

Acknowledgments

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No [694980] SYNTH: Synthesising Inductive Data Models). OK was supported by Czech Science Foundation project “Generative Relational Models” (20-19104Y) and partially by the OP VVV project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”. Part of this work was done while NK was visiting CTU in Prague, supported by Research Center for Informatics.

Appendix A. Missing Proofs

A.1 Proof of Lemma 1

In this section, we present the detailed proof of Lemma 1.

Proof. Let us denote the variables in \mathbf{Z} that are marked on the top (requisite) by \mathbf{Z}_\star and that are not marked on the top (not requisite) by $\mathbf{Z}_\bar{\star}$. The required probability μ is then given by,

$$\mu = P(\mathbf{x}_q \mid \mathbf{e}) = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star, \mathbf{z}_\bar{\star}} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{z}_\bar{\star}, \mathbf{e}) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star, \mathbf{z}_\bar{\star}} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{z}_\bar{\star}, \mathbf{e})} = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}) f(\mathbf{x}) \sum_{\mathbf{z}_\bar{\star}} P(\mathbf{z}_\bar{\star} \mid \mathbf{x}, \mathbf{z}_\star, \mathbf{e})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}) \sum_{\mathbf{z}_\bar{\star}} P(\mathbf{z}_\bar{\star} \mid \mathbf{x}, \mathbf{z}_\star, \mathbf{e})}$$

Since $\sum_{\mathbf{z}_\bar{\star}} P(\mathbf{z}_\bar{\star} \mid \mathbf{x}, \mathbf{z}_\star, \mathbf{e}) = 1$, we can write,

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e})}$$

Now let us denote the observed variables in \mathbf{E} that are visited (requisite) by \mathbf{E}_r and those that are not visited (not requisite) by \mathbf{E}_n . We can write,

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r) P(\mathbf{e}_n | \mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r) P(\mathbf{e}_n | \mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r)}$$

The variables in $\mathbf{X} \cup \mathbf{Z}_\star$ pass the Bayes-balls to all their parents and all their children, but \mathbf{E}_n is not visited by these balls. The correctness of the Bayes-ball algorithm ensures that there is no active path from $\mathbf{X} \cup \mathbf{Z}_\star$ to any E_n in \mathbf{E}_n given \mathbf{E}_r . Thus $\mathbf{X}, \mathbf{Z}_\star \perp \mathbf{E}_n | \mathbf{E}_r$ and $P(\mathbf{e}_n | \mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r) = P(\mathbf{e}_n | \mathbf{e}_r)$. After cancelling out the common term $P(\mathbf{e}_n | \mathbf{e}_r)$, we get,

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_r)}$$

Now let us denote observed variables in \mathbf{E}_r that are only visited by \mathbf{E}_\star and that are visited as well as marked on top by \mathbf{E}_\star . After canceling out the common term, we get the desired result,

$$\mu = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star | \mathbf{e}_\star) P(\mathbf{e}_\star) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star | \mathbf{e}_\star) P(\mathbf{e}_\star)} = \frac{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star | \mathbf{e}_\star) f(\mathbf{x})}{\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star | \mathbf{e}_\star)}$$

□

A.2 Proof of Proposition 2

In this section, we present the detailed proof of Proposition 2.

Proof. First, we show that if $\mathbf{rv}(\mathbf{A})$ is in the least Herbrand model of $rv(\mathbb{P})$ (denoted by $M_{rv(\mathbb{P})}$) then \mathbb{P} defines \mathbf{A} as an RV: Since \mathbb{P} is well-defined, there will be at least one clause for \mathbf{A} in $ground(\mathbb{P})_\omega$ whose body is true with respect to ω due to the exhaustiveness condition. So, \mathbb{P} defines \mathbf{A} as an RV.

Now, we show the converse: Let ω be an assignment where each \mathbf{B} is assigned a value if $\mathbf{rv}(\mathbf{B})$ is in $M_{rv(\mathbb{P})}$. Given ω , we construct a ground program $ground(\mathbb{P})_\omega$ such that there is a clause $\mathbf{A} \sim \mathbf{D} \leftarrow \text{Body}$ in $ground(\mathbb{P})_\omega$ with atoms $\{\mathbf{B}_1 \cong \mathbf{V}_1, \dots, \mathbf{B}_m \cong \mathbf{V}_m\}$ in Body and these atoms are true with respect to ω . Then, due to the way DC# programs have been defined in Definition 5, there must be a definite clause $\mathbf{rv}(\mathbf{A}) \leftarrow \mathbf{rv}(\mathbf{B}_1), \dots, \mathbf{rv}(\mathbf{B}_m)$ in $ground(rv(\mathbb{P}))$. Now, $\{\mathbf{rv}(\mathbf{B}_1), \dots, \mathbf{rv}(\mathbf{B}_m)\}$ are in $M_{rv(\mathbb{P})}$ since those RVs are assigned in ω . So, $\mathbf{rv}(\mathbf{A})$ will be in $M_{rv(\mathbb{P})}$ since the body of the definite clause will be true. This completes the proof. □

A.3 Proof of Proposition 3

In this section, we present the detailed proof of Proposition 3.

Proof. First, we show that if \mathbf{A} directly influences \mathbf{B} then $\text{pa}(\mathbf{B}, \mathbf{A})$ is in the least Herbrand model of $dep(\mathbb{P})$: Since \mathbf{A} directly influences \mathbf{B} there must be a clause $\mathcal{C} \equiv \mathbf{B} \sim \mathbf{D} \leftarrow \mathbf{B}_1, \dots, \mathbf{B}_n \in ground(\mathbb{P})_\omega$ for some possible world ω such that $n > 0$, \mathbf{A} is a RV term in the body of \mathcal{C} , and the body is true with respect to ω . This implies that there must be a clause $\mathbf{rv}(\mathbf{B}) \leftarrow \mathbf{rv}(\mathbf{A}_1), \dots, \mathbf{rv}(\mathbf{A}_m)$ in $ground(rv(\mathbb{P}))$ such that $\{\mathbf{A}_1, \dots, \mathbf{A}_m\}$ is the set of RV terms

in the body of \mathcal{C} and $\mathbf{A} \in \{\mathbf{A}_1, \dots, \mathbf{A}_m\}$. Consequently, there must be a definite clause $\mathcal{C}' \equiv \text{pa}(\mathbf{B}, \mathbf{A}) \leftarrow \text{rv}(\mathbf{B}), \text{rv}(\mathbf{A}_1), \dots, \text{rv}(\mathbf{A}_m)$ in $\text{ground}(\text{dep}(\mathbb{P}))$ by construction. $\{\mathbf{A}_1, \dots, \mathbf{A}_m\}$ is assigned in ω so due to Proposition 2: $\text{rv}(\mathbb{P}) \models (\text{rv}(\mathbf{A}_1), \dots, \text{rv}(\mathbf{A}_m))$, consequently, $\text{rv}(\mathbb{P}) \models \text{rv}(\mathbf{B})$. This is also true in $\text{dep}(\mathbb{P})$ since all clauses in $\text{rv}(\mathbb{P})$ are in $\text{dep}(\mathbb{P})$. So, the body of clause \mathcal{C}' will be true and $\text{dep}(\mathbb{P}) \models \text{pa}(\mathbf{B}, \mathbf{A})$.

Now, we show the converse: Since $\text{dep}(\mathbb{P}) \models \text{pa}(\mathbf{B}, \mathbf{A})$, there must be a definite clause $\text{pa}(\mathbf{B}, \mathbf{A}) \leftarrow \text{rv}(\mathbf{B}), \text{rv}(\mathbf{A}_1), \dots, \text{rv}(\mathbf{A}_m)$ in $\text{ground}(\text{dep}(\mathbb{P}))$ such that $\mathbf{A} \in \{\mathbf{A}_1, \dots, \mathbf{A}_m\}$ and the clause's body is true in $\text{ground}(\text{dep}(\mathbb{P}))$. This implies that there must be a definite clause $\text{rv}(\mathbf{B}) \leftarrow \text{rv}(\mathbf{A}_1), \dots, \text{rv}(\mathbf{A}_m) \in \text{ground}(\text{rv}(\mathbb{P}))$ whose body is true in $\text{ground}(\text{rv}(\mathbb{P}))$. Due to the way DC# programs are defined (Definition 5), there must be a distributional clause $\mathbf{B} \sim \mathbf{D} \leftarrow \text{Body} \in \text{ground}(\mathbb{P})_\omega$ for some ω such that $\{\mathbf{A}_1 \cong \mathbf{V}_1, \dots, \mathbf{A}_m \cong \mathbf{V}_m\}$ belongs to Body that is true in ω . So, \mathbf{A} directly influences \mathbf{B} , which completes the proof. \square

A.4 Proof of Theorem 5

In this section, we present the detailed proof of Theorem 5.

Proof. The expectation $\mathbb{E}_{Q_\star}[W_{\dot{\mathbf{e}}_\star}]$ is given by

$$\sum_{\mathbf{x}, \mathbf{z}_\star} \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\star} P(u_i \mid \mathbf{pa}(U_i)) \prod_{v_i \in \dot{\mathbf{e}}_\star} P(v_i \mid \mathbf{pa}(V_i)).$$

The basis $\dot{\mathbf{S}}_\star$ is a subset of $\mathbf{X} \cup \mathbf{Z}_\star$ by Definition 16. Let us denote $(\mathbf{X} \cup \mathbf{Z}_\star) \setminus \dot{\mathbf{S}}_\star$ by \mathbf{Z}_\diamond . We can now rewrite the expectation as follows,

$$\sum_{\dot{\mathbf{S}}_\star, \mathbf{z}_\diamond} \prod_{u_i \in \dot{\mathbf{e}}_\star \cup \dot{\mathbf{S}}_\star} P(u_i \mid \mathbf{pa}(U_i)) \prod_{v_i \in \mathbf{z}_\diamond} P(v_i \mid \mathbf{pa}(V_i)).$$

We will show that $Pa \notin \mathbf{Z}_\diamond$ for any $Pa \in \mathbf{Pa}(U_i)$, which will then allow us to push the summation over \mathbf{z}_\diamond inside. Let us consider two cases:

- For $U_i \in \dot{\mathbf{E}}_\star$, let $Pa \in \mathbf{Pa}(U_i)$ be an unobserved parent of U_i , then there will be a direct causal trail from Pa to U_i , consequently Pa will be in the set $\dot{\mathbf{S}}_\star$.
- For $U_i \in \dot{\mathbf{S}}_\star$, there will be a causal trail $U_i \rightarrow \dots B_j \dots \rightarrow E$ such that $E \in \dot{\mathbf{E}}_\star$ and such that either no B_i is observed or there is no B_i . Let $Pa \in \mathbf{Pa}(U_i)$ be an unobserved parent of U_i then there will be a direct causal trail from Pa to U_i , consequently, there will be such causal trail from Pa to E and Pa will be in the set $\dot{\mathbf{S}}_\star$.

Hence, we push the summation over \mathbf{z}_\diamond inside and use the fact that $\sum_{\mathbf{z}_\diamond} \prod_{v_i \in \mathbf{z}_\diamond} P(v_i \mid \mathbf{pa}(V_i)) = 1$, to get the desired result. \square

A.5 Proof of Theorem 6

In this section, we present the detailed proof of Theorem 6.

Proof. Since $\mathbf{X}, \mathbf{Z}_\star, \mathbf{E}_\star, \mathbf{E}_\star$ are variables of the Bayesian network \mathcal{B} and they form a sub-network \mathcal{B}_\star such that \mathbf{E}_\star do not have any parent, we can always write,

$$P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star) = \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\star \cup \mathbf{e}_\star} P(u_i \mid \mathbf{pa}(U_i)) \prod_{v_i \in \mathbf{z}_\star \cup \mathbf{e}_\star} P(v_i \mid \mathbf{pa}(V_i))$$

such that $p \in \mathbf{x} \cup \mathbf{z}_\star \cup \mathbf{e}_\star \cup \mathbf{e}_\star$ for all $p \in \mathbf{pa}(U_i)$ or $p \in \mathbf{pa}(V_i)$. Now consider the summation over all possible assignments of variables in $\mathbf{X}, \mathbf{Z}_\star$, that is: $\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star)$. We will write this summation as a summation over contextual assignments $\psi \in \Psi$. The collection of the ranges of all $\psi \in \Psi$ will have all possible assignments of variables in $\mathbf{X}, \mathbf{Z}_\star$ without any duplicates as assignments $\psi \in \Psi$ are mutually exclusive. So, we can always write,

$$\sum_{\mathbf{x}, \mathbf{z}_\star} P(\mathbf{x}, \mathbf{z}_\star, \mathbf{e}_\star \mid \mathbf{e}_\star) = \sum_{\psi \in \Psi} \sum_{\mathbf{z}_\dagger[\psi]} P(\mathbf{x}[\psi], \mathbf{z}_\dagger[\psi], \mathbf{z}_\dagger[\psi], \mathbf{e}_\dagger[\psi], \mathbf{e}_\dagger[\psi] \mid \mathbf{e}_\star) \quad (8)$$

In the above equation, notice that the inner summation is over the range of ψ . To simplify notation, from now we denote $\{\mathbf{x}[\psi], \mathbf{z}_\dagger[\psi], \mathbf{Z}_\dagger[\psi], \mathbf{e}_\dagger[\psi], \mathbf{e}_\dagger[\psi]\}$ by $\{\mathbf{x}, \mathbf{z}_\dagger, \mathbf{Z}_\dagger, \mathbf{e}_\dagger, \mathbf{e}_\dagger\}$. After using the definition of contextual assignments, we have that,

$$P(\mathbf{x}, \mathbf{z}_\dagger, \mathbf{z}_\dagger, \mathbf{e}_\dagger, \mathbf{e}_\dagger \mid \mathbf{e}_\star) = \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(u_i \mid \mathbf{ppa}(U_i)) \prod_{v_i \in \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(v_i \mid \mathbf{pa}(V_i))$$

Since $p \notin \mathbf{z}_\dagger$ for any $p \in \mathbf{ppa}(U_i)$, we can push the summation over \mathbf{z}_\dagger inside to get,

$$\sum_{\psi \in \Psi} \sum_{\mathbf{z}_\dagger} P(\mathbf{x}, \mathbf{z}_\dagger, \mathbf{z}_\dagger, \mathbf{e}_\dagger, \mathbf{e}_\dagger \mid \mathbf{e}_\star) = \sum_{\psi \in \Psi} \prod_{u_i \in \mathbf{x} \cup \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(u_i \mid \mathbf{ppa}(U_i)) \sum_{\mathbf{z}_\dagger} \prod_{v_i \in \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(v_i \mid \mathbf{pa}(V_i)). \quad (9)$$

However, we get a strange term $\sum_{\mathbf{z}_\dagger} \prod_{v_i \in \mathbf{z}_\dagger \cup \mathbf{e}_\dagger} P(v_i \mid \mathbf{pa}(V_i))$. Let \mathbf{S}_\dagger denote the basis of residual \mathbf{e}_\dagger . We have that $\mathbf{S}_\dagger \subseteq \mathbf{Z}_\dagger$ by Definition 17. Let us denote $\mathbf{Z}_\dagger \setminus \mathbf{S}_\dagger$ with \mathbf{Z}_\diamond . Now the strange term can be rewritten as,

$$\sum_{\mathbf{S}_\dagger, \mathbf{Z}_\diamond} \prod_{u_i \in \mathbf{e}_\dagger \cup \mathbf{S}_\dagger} P(u_i \mid \mathbf{pa}(U_i)) \prod_{v_i \in \mathbf{Z}_\diamond} P(v_i \mid \mathbf{pa}(V_i)).$$

In the proof of Theorem 5, we showed that the summation over variables not in \mathbf{S}_\dagger can be pushed inside; hence, \mathbf{Z}_\diamond can be pushed inside. After using the fact that $\sum_{\mathbf{z}_\diamond} \prod_{v_i \in \mathbf{z}_\diamond} P(v_i \mid \mathbf{pa}(V_i)) = 1$, we conclude that the strange term is actually the expectation $\mathbb{E}_{Q_\star} [W_{\mathbf{e}_\dagger}]$. Using Equation (3), (8), (9) and rearranging terms, the result follows. \square

A.6 Proof of Lemma 7

In this section, we present the detailed proof of Lemma 7.

Proof. It is clear that a subset of unobserved variables is assigned. Let \mathbf{Z}_\dagger be a set of unobserved variables left unassigned. Let $E \in \mathbf{E}_\star$ be an observed variable. Consider two cases:

- All ancestors of E are in $\mathbf{Z}_\dagger \cup \mathbf{E}_\star \cup \mathbf{E}_\star$.

- Some ancestors of E are in $\mathbf{Z}_{\dagger} \cup \mathbf{E}_{\star} \cup \mathbf{E}_{\star}$ and some are in $\mathbf{X} \cup \mathbf{Z}_{\dagger}$. Let $A \in \mathbf{X} \cup \mathbf{Z}_{\dagger}$ and let $A \rightarrow \dots B_i \dots \rightarrow E$ be a causal trail. Some B_i are observed in all such trails.

Clearly, E will not be visited from any parent in the first case, and in the second case, the visit will be blocked by observed variables. Consequently, E will not be weighted, which completes the proof. \square

A.7 Proof of Theorem 8

In this section, we present the detailed proof of Theorem 8.

Proof. Variables in \mathbf{Z}_{\dagger} are not assigned in the simulation; hence, it follows immediately from Lemma 9 that the assignment is contextual. Assume by contradiction that $A \in \mathbf{X} \cup \mathbf{Z}_{\dagger}$, $E \in \mathbf{E}_{\dagger}$ and there is a causal trail $A \rightarrow \dots B_i \dots \rightarrow E$ such that no B_i is observed or there is no B_i . Since A is assigned, all children of A will be visited, and following the trail, the variable E will also be visited from its parent since there is no observed variable in the trail to block the visit. Consequently, E will be weighted, which contradicts our assumption that E is not weighted. Hence, the assignment is also safe. \square

A.8 Proof of Lemma 9

In this section, we present the detailed proof of Lemma 9.

Proof. Since A is assigned/weighted and rules in \mathbb{P} are exhaustive, a rule $\mathcal{R} \in \mathbb{P}$ with A in its head must have fired. Let \mathbf{d} be a body and \mathcal{D} be a distribution in the head of \mathcal{R} . Since each $d_i \in \mathbf{d}$ must be true for \mathcal{R} to fire, $\mathbf{d} \subseteq \mathbf{c}$. We assume that rules in \mathbb{P} are mutually exclusive. Thus, among all rules for A , only \mathcal{R} will fire even when an assignment of some variables in \mathbf{B} is also given. Hence, by definition of the rule \mathcal{R} , we have that,

$$\mathcal{D} = P(A \mid \mathbf{d}) = P(A \mid \mathbf{c}) = P(A \mid \mathbf{c}, \mathbf{B})$$

\square

A.9 Proof of Theorem 10

In this section, we present the detailed proof of Theorem 10.

Proof. Since random variable A is sampled/weighted, k clauses out of n ground distributional clauses for A must have fired. Those $n - k$ clauses that do not fire must have at least one atom in their body that is false. Let \mathbf{d} be an assignment of random variables in the body of the k clauses and \mathcal{D} be a distribution obtained after applying a combining rule on the multiset of distributions in the head of the k clauses. Since each $d \in \mathbf{d}$ must be true for k clauses to fire, $\mathbf{d} \subseteq \mathbf{c}$. Even if an assignment of some variables in \mathbf{B} is also given, the $n - k$ clauses cannot fire since one atom in the body of these clauses is already false. Hence, by definition, we have that

$$\mathcal{D} = P(A \mid \mathbf{d}) = P(A \mid \mathbf{c}) = P(A \mid \mathbf{c}, \mathbf{B})$$

\square

Appendix B. Advanced Constructs in DC# Framework

This section describes advanced constructs such as negation, aggregates, and statistical models in the DC# framework, which are useful while writing complex relational models (Kumar et al., 2021).

Example 18. *The following program illustrates advanced constructs such as negation, aggregates, and statistical models.*

```

client(ann) ~ val(true).
loan(l.1) ~ bernoulli(0.9).
loan(l.2) ~ bernoulli(0.9).
age(C) ~ gaussian(40, 10.5) ← client(C) ≅ true.
has_loan(C, L) ~ bernoulli(0.2) ← client(C) ≅ true, loan(L) ≅ true.
status(L) ~ discrete([0.3 : appr, 0.7 : decl]) ← loan(L) ≅ true.
credit_score(C) ~ gaussian(700, 10.9) ← has_loan(C, L) ≅ true, status(L) ≅ appr.
credit_score(C) ~ gaussian(600, 20.5) ← has_loan(C, L) ≅ true, status(L) ≅ decl.
credit_score(C) ~ gaussian(750, 30.2) ← has_loan(C, L) ≅ true, \+status(L) ≅ ..
credit_score(C) ~ gaussian(M, 15.9) ← age(C) ≅ Y,
    mode(X, (has_loan(C, L) ≅ true, status(L) ≅ X), appr), linear([Y], [20.1, 30.9], M).

```

Next, we explain the advanced constructs of the above program one by one in detail.

Negation The exhaustiveness condition for well-defined programs is somewhat contrived. In practice, it is difficult to specify conditional probability distributions of a RV given all possible assignments of its parents because real-world data is often inadequate and contains missing values. In such cases, it is convenient to describe models using negative literals in the bodies of DCs.

The second last clause in the program of Example 18 has a negation in the body. The negation is interpreted as *negation as failure* as usual in logic programming. We understand from the above program that the loan’s status can take three values: `appr` (“approved”), `decl` (“declined”), and “undefined”. That is, its value can be unknown (or undefined) in some possible worlds. The loan’s status takes the undefined value with a probability of 1 when the loan’s identity is false, and the program specifies a probability distribution over such worlds.

To avoid *floundering* (Nilsson & Małuszyński, 1995), we allow writing only those clauses $A_0 \sim D \leftarrow L_1, \dots, L_n$ in DC# programs, which satisfy this condition: if a logical variable J occurs in the RV term of a negative literal L_i then J occurs in the RV term of a positive literal in $\{A_0, L_1, \dots, L_{i-1}\}$. This means writing *unsafe negations* are not allowed. For example, writing the clause

$$\text{credit_score}(C) \sim \text{gaussian}(500, 30.2) \leftarrow \backslash + \text{status}(L) \cong ..$$

is not allowed because the logical variable L occurs in the RV term of the negative literal but it does not occur in any positive literal.

Aggregates An alternative approach to combining rules, extensively used for describing models for relational data, is the aggregation function or aggregate. It is a function that maps a multiset of values to a single value. Standard examples are average (if values are numerical), mode (most frequently occurring value), maximum, minimum, cardinality, etc. For example, the mode maps a multiset $\{\text{appr}, \text{decl}, \text{appr}\}$ to a value appr . To express these functions, built-in aggregation predicates are used in conjunction with the second-order predicate `findall/3` as follows: `findall(T,G,L),aggr(L,R)`, where T is a target variable such that it occurs in a goal G , and L unifies with the multiset of the instantiations that T gets successively on backtracking over G . This is just like Prolog’s `findall/3` predicate (Wielemaker, et al., 2011). The atom `findall(T,G,L)` succeeds with an empty multiset if goal G has no solutions. A built-in predicate `aggr/2` implements an aggregation function (e.g., mode) that maps the multiset L to a single value R . When the multiset L is empty, the aggregation atom `aggr(L,R)` fails. We will use a shorthand notation `aggr(T,G,R)` for `findall(T,G,L),aggr(L,R)` and call it an *aggregation atom*.

The last clause in the program of Example 18 has an aggregation atom in its body. Here, the idea is to consider the joint influence of statuses of all loans that a client has on the distribution of the client’s credit score (instead of letting each loan have its own independent influence as in Example 7). There are three situations in which the aggregation atom will fail, and the body of the last clause will be true: (i) the client has no loan, (ii) it is undefined that the client has loans or not, (iii) the client has loans, but statuses of those loans are undefined.

The support for aggregates, which in turn requires support for the second-order predicate `findall/3`, takes `DC#` outside the domain of first-order probabilistic logic. However, the semantics of programs with aggregates can also be understood in terms of the ground programs. Just like Prolog, atom `aggr(T,G,R)` gets its truth value after checking for truths of all instantiations of goal G in the ground program.

Statistical Model Atom It maps outcomes of RVs in the body of a DC to parameters of the distribution in the head. Formally, a DC with a statistical model is a clause of the form $A_0 \sim D_\phi \leftarrow A_1, \dots, A_n, M_\psi$, where M_ψ is an atom implementing a mathematical function that relates the values of RVs in $\{A_1, \dots, A_n\}$ with parameters ϕ in distribution D_ϕ via parameters ψ .

An example is shown in the last clause of the program of Example 18. Here, `linear` is a built-in predicate that implements a linear model relating the value of client’s age and mean M of the distribution in the head using parameters $\psi = [20.1, 30.9]$ as follows: `M is 20.1 × X + 30.9`.

Many well-known statistical models, such as linear regression, logistic regression, soft-max regression, etc., can be used to describe very complex relational models (Kumar et al., 2021). Importantly, we have integrated the full expressiveness of logic programming with the strengths of statistical models in learning intricate patterns.

B.1 Inference in Programs with Advanced Constructs

For inference, first, we should identify RVs from `DC#` programs with advanced constructs.

Definition 19. Let \mathbb{P} be a well-defined DC# program with advanced constructs. The RV set $rv(\mathbb{P})$ of the program is the set of definite clauses obtained by transforming each clause $A \sim D \leftarrow L_1, \dots, L_m \in \mathbb{P}$ like this:

1. Let Body be the empty set. For each positive or negative literal L_i of the form $T \cong V$, an atom $rv(T)$ is added to Body .
2. A clause $rv(A) \leftarrow \text{Body}$ is added to $rv(\mathbb{P})$.

Example 19. The RV set for the program in Example 18 is:

```

rv(client(ann)).
rv(loan(1.1)).
rv(loan(1.2)).
rv(age(C)) ← rv(client(C)).
rv(has_loan(C,L)) ← rv(client(C)),rv(loan(L)).
rv(status(L)) ← rv(loan(L)).
rv(credit_score(C)) ← rv(has_loan(C,L)),rv(status(L)).
rv(credit_score(C)) ← rv(age(C)).
    
```

The above definition differs from Definition 6 in only one way. It additionally specifies the way negative literals of the form $\setminus +T \cong V$ should be handled. However, nothing special is done for them because just like positive literals that compare outcomes of ground RV terms with values, negative literals compare the outcomes with values or with the “undefined” value. Additionally, comparison and statistical model atoms are ignored while constructing the RV set. This is because they do not contain RV terms. The aggregation atoms do contain RV terms, but they do not constrain instantiations of logical variables appearing in other atoms. So, they are also ignored.

However, we must ensure that Proposition 2 is valid even after using negations that relax the exhaustiveness condition. Now, it may happen that $rv(\mathbb{P}) \models rv(A)$ but no distribution is specified for A in some ground programs. In this case, however, the “undefined” value is assigned to A , so the exhaustiveness condition is still implicitly present and Proposition 2 is valid.

Next, we should be able to identify direct influence relationships among RVs. However, the second-order aggregation atoms complicate the construction of dependency sets. We add an additional rule in Definition 8 to deal with it.

Definition 20. Let $rv(\mathbb{P})$ be the RV set of a well-defined DC# program \mathbb{P} with advanced constructs. The dependency set $dep(\mathbb{P})$ is the union of $rv(\mathbb{P})$ and the set of definite clauses $pa(\mathbb{P})$ obtained by transforming each clause $C \equiv A \sim D \leftarrow L_1, \dots, L_m \in \mathbb{P}$ with non empty body as follows:

1. Let $rv(A) \leftarrow rv(T_1), \dots, rv(T_n)$ be the clause in $rv(\mathbb{P})$ corresponding to C .
2. For each $rv(T_i)$, a clause $pa(A, T_i) \leftarrow rv(A), rv(T_1), \dots, rv(T_n)$ is added to $pa(\mathbb{P})$.
3. For each literal L_i of the form $\text{aggr}(T, Q, R)$,
 - (a) Let Q be of the form (Q_1, \dots, Q_n) and let Body be the empty set.

(b) For each positive or negative literal Q_i of the form $T \cong V$, an atom $rv(T)$ is added to Body.

(c) For each positive or negative literal Q_i of the form $T \cong V$, a clause

$$pa(A, T) \leftarrow rv(A), rv(T_1), \dots, rv(T_n), Body$$

is added to $pa(\mathbb{P})$.

However, note that in Prolog, all free variables appearing in `findall/3` are bound with the existential operator. So, all free variables in `aggr/3` should be renamed before constructing the dependency set.

Example 20. In addition to the RV set of Example 19, the dependency set of Example 18 contains the following clauses. Notice the renamings in the last two clauses.

```
pa(age(C), client(C)) ← rv(age(C)), rv(client(C)).
pa(has_loan(C, L), client(C)) ← rv(has_loan(C, L)), rv(client(C)), rv(loan(L)).
pa(has_loan(C, L), loan(L)) ← rv(has_loan(C, L)), rv(client(C)), rv(loan(L)).
pa(status(L), loan(L)) ← rv(status(L)), rv(loan(L)).
pa(credit_score(C), has_loan(C, L)) ← rv(credit_score(C)), rv(has_loan(C, L)), rv(status(L)).
pa(credit_score(C), status(L)) ← rv(credit_score(C)), rv(has_loan(C, L)), rv(status(L)).
pa(credit_score(C), age(C)) ← rv(credit_score(C)), rv(age(C)).
pa(credit_score(C), has_loan(C, L1)) ← rv(credit_score(C)), rv(age(C)),
    rv(has_loan(C, L1)), rv(status(L1)).
pa(credit_score(C), status(L2)) ← rv(credit_score(C)), rv(age(C)),
    rv(has_loan(C, L2)), rv(status(L2)).
```

After identifying RVs and direct influences among them, the simulation of programs defined in Algorithm 5 is straightforward. One can easily add rules to the proof procedure (Algorithm 6) for resolving these advanced constructs. In the implementation, predicates `rv/1`, `pa/2` should be *tabled* (Warren, 1992) for efficiency.

References

- Ahmadi, B., Kersting, K., Mladenov, M., & Natarajan, S. (2013). Exploiting symmetries for scaling loopy belief propagation and relational training. *Machine Learning*, 92, 91–132.
- Alberti, M., Bellodi, E., Cota, G., Riguzzi, F., & Zese, R. (2017). cplint on swish: Probabilistic logical inference with a web browser. *Intelligenza Artificiale*, 11(1), 47–64.
- Beame, P., Van den Broeck, G., Gribkoff, E., & Suciu, D. (2015). Symmetric weighted first-order model counting. In *Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pp. 313–328.
- Belle, V., Passerini, A., & Van Den Broeck, G. (2015). Probabilistic inference in hybrid domains by weighted model integration. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, p. 2770–2776.

- Boutilier, C., Friedman, N., Goldszmidt, M., & Koller, D. (1996). Context-specific independence in bayesian networks. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pp. 115–123. Morgan Kaufmann Publishers Inc.
- Chavira, M., & Darwiche, A. (2008). On probabilistic inference by weighted model counting. *Artificial Intelligence*, 172(6-7), 772–799.
- Chen, Y., Yang, Y., Natarajan, S., & Ruoizzi, N. (2020). Lifted hybrid variational inference. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pp. 4237–4244.
- Choi, J., Braz, R. d. S., & Bui, H. H. (2011). Efficient methods for lifted inference with aggregate factors. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI’11*, p. 1030–1036. AAAI Press.
- Corander, J., Hyttinen, A., Kontinen, J., Pensar, J., & Väänänen, J. (2019). A logical approach to context-specific independence. *Annals of Pure and Applied Logic*, 170(9), 975–992.
- Darwiche, A., & Marquis, P. (2002). A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17, 229–264.
- De Raedt, L., Kersting, K., Natarajan, S., & Poole, D. (2016). Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 10(2), 1–189.
- De Raedt, L., & Kimmig, A. (2015). Probabilistic (logic) programming concepts. *Machine Learning*, 100(1), 5–47.
- De Raedt, L., Kimmig, A., & Toivonen, H. (2007). Problog: A probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI’07*, p. 2468–2473.
- De Salvo Braz, R., Amir, E., & Roth, D. (2006). Mpe and partial inversion in lifted probabilistic variable elimination. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 2, AAAI’06*, p. 1123–1130.
- de Salvo Braz, R., Natarajan, S., Bui, H., Shavlik, J., & Russell, S. (2009). Anytime lifted belief propagation. *Proc. SRL*, 9.
- Elidan, G. (2001). Bayesian Network Repository.. <https://www.cse.huji.ac.il/galel/Repository/>.
- Feldstein, J., & Belle, V. (2021). Lifted reasoning meets weighted model integration. In *37th Conference on Uncertainty in Artificial Intelligence*.
- Fierens, D., Van den Broeck, G., Renkens, J., Shterionov, D., Gutmann, B., Thon, I., Janssens, G., & De Raedt, L. (2015). Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory and Practice of Logic Programming*, 15(3), 358–401.
- Friedman, N., Getoor, L., Koller, D., & Pfeffer, A. (1999). Learning probabilistic relational models. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI’99*, p. 1300–1307.

- Friedman, T., & Van Den Broeck, G. (2018). Approximate knowledge compilation by online collapsed importance sampling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, p. 8035–8045.
- Getoor, L., Friedman, N., Koller, D., Pfeffer, A., & Taskar, B. (2007). Probabilistic relational models. *Introduction to statistical relational learning*, 8.
- Getoor, L., & Grant, J. (2006). Prl: A probabilistic relational language. *Machine Learning*, 62(1), 7–31.
- Getoor, L., & Taskar, B. (2007). Lifted first-order probabilistic inference. In *Statistical Relational Learning*.
- Gogate, V., & Domingos, P. M. (2011). Probabilistic theorem proving. *Communications of the ACM*, 59, 107 – 115.
- Gutmann, B., Jaeger, M., & De Raedt, L. (2010). Extending problog with continuous distributions. In *International Conference on Inductive Logic Programming*, Vol. 6489 of *Lecture Notes in Computer Science*, pp. 76–91. Springer.
- Gutmann, B., Thon, I., Kimmig, A., Bruynooghe, M., & De Raedt, L. (2011). The magic of logical inference in probabilistic programming. *Theory and Practice of Logic Programming*, 11(4-5), 663–680.
- Heckerman, D., Meek, C., & Koller, D. (2004). Probabilistic models for relational data. Tech. rep., Technical Report MSR-TR-2004-30, Microsoft Research.
- Islam, M. A., Ramakrishnan, C., & Ramakrishnan, I. (2012). Inference in probabilistic logic programs with continuous random variables. *Theory and Practice of Logic Programming*, 12(4-5), 505–523.
- Jaeger, M. (2007). Parameter learning for relational bayesian networks. In *Proceedings of the 24th international conference on Machine learning*, pp. 369–376.
- Kazemi, S. M., Kimmig, A., Van den Broeck, G., & Poole, D. (2016). New liftable classes for first-order probabilistic inference. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, p. 3125–3133.
- Kersting, K., & De Raedt, L. (2001). Adaptive bayesian logic programs. In *Proceedings of the 11th International Conference on Inductive Logic Programming*, ILP '01, p. 104–117, Berlin, Heidelberg. Springer-Verlag.
- Kersting, K., & De Raedt, L. (2007). Bayesian logic programming: Theory and tool. In *Introduction to Statistical Relational Learning*. MIT Press.
- Kisynski, J., & Poole, D. (2009). Lifted aggregation in directed first-order probabilistic models.. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI'09, pp. 1922–1929.
- Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- Koller, D., Friedman, N., Džeroski, S., Sutton, C., McCallum, A., Pfeffer, A., Abbeel, P., Wong, M.-F., Heckerman, D., Meek, C., et al. (2007). *Introduction to statistical relational learning*. MIT press.

- Kowalski, R. (1974). Predicate logic as programming language. In *IFIP congress*, Vol. 74, pp. 569–544.
- Kumar, N., & Kuželka, O. (2021). Context-specific likelihood weighting. In *International Conference on Artificial Intelligence and Statistics*, pp. 2125–2133. PMLR.
- Kumar, N., Kuželka, O., & De Raedt, L. (2021). Learning distributional programs for relational autocompletion. *Theory and Practice of Logic Programming*, 21, 1–34.
- Li, L., & Russell, S. J. (2013). The blog language reference. Tech. rep., Technical Report UCB/EECS-2013-51, EECS Department, University of California.
- Mateescu, R., & Dechter, R. (2009). Mixed deterministic and probabilistic networks. *Annals of Mathematics and Artificial Intelligence*, 54, 3–51.
- Michels, S., Hommersom, A., & Lucas, P. J. F. (2016). Approximate probabilistic inference with bounded error for hybrid probabilistic logic programming. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, p. 3616–3622. AAAI Press.
- Milch, B., Marthi, B., Sontag, D., Russell, S., Ong, D. L., & Kolobov, A. (2005). Approximate inference for infinite contingent bayesian networks. In *International Workshop on Artificial Intelligence and Statistics, AISTATS'05*, pp. 238–245. PMLR.
- Milch, B., & Russell, S. J. (2009). Extending bayesian networks to the open-universe case. In *R. Dechter, H. Geffner, and J. Y. Halpern, eds. Heuristics, Probability and Causality: A Tribute to Judea Pearl*.
- Minka, T., Winn, J., Guiver, J., & Knowles, D. (2010). Infer .net 2.4, 2010. microsoft research cambridge..
- Natarajan, S., Tadepalli, P., Dietterich, T. G., & Fern, A. (2008). Learning first-order probabilistic models with combining rules. *Annals of Mathematics and Artificial Intelligence*, 54(1), 223–256.
- Ngo, L., & Haddawy, P. (1995). Probabilistic logic programming and bayesian networks. In *Asian computing science conference*, pp. 286–300. Springer.
- Ngo, L., & Haddawy, P. (1997). Answering queries from context-sensitive probabilistic knowledge bases. *Theoretical Computer Science*, 171(1-2), 147–177.
- Niepert, M. (2012). *Lifted Probabilistic Inference: An MCMC Perspective*. Universitätsbibliothek Mannheim.
- Nilsson, U., & Małuszyński, J. (1995). *Logic, programming and Prolog*. Wiley Chichester.
- Nitti, D., De Laet, T., & De Raedt, L. (2016). Probabilistic logic programming for hybrid relational domains. *Machine Learning*, 103(3), 407–449.
- Poole, D. (1993). Probabilistic horn abduction and bayesian networks. *Artificial intelligence*, 64(1), 81–129.
- Poole, D. (1997). Probabilistic partial evaluation: Exploiting rule structure in probabilistic inference. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'97*, p. 1284–1291.

- Poole, D. (2003). First-order probabilistic inference. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI'03*, p. 985–991.
- Poole, D. (2008). The independent choice logic and beyond. In *Probabilistic inductive logic programming*, pp. 222–243. Springer.
- Poole, D., & Zhang, N. L. (2003). Exploiting contextual independence in probabilistic inference. *Journal of Artificial Intelligence Research*, 18, 263–313.
- Poole, D. L., & Mackworth, A. K. (2010). *Artificial Intelligence: foundations of computational agents*. Cambridge University Press.
- Sato, T., & Kameya, Y. (1997). Prism: A language for symbolic-statistical modeling. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'97*, p. 1330–1335.
- Schulte, O., & Routley, K. (2014). Aggregating predictions vs. aggregating features for relational classification. In *Computational Intelligence and Data Mining (CIDM), 2014 IEEE Symposium on*, pp. 121–128. IEEE. <https://relational.fit.cvut.cz/dataset/NBA>.
- Shachter, R. (1998). Bayes-ball: The rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, UAI'97*, p. 480–487.
- Shen, Y., Choi, A., & Darwiche, A. (2016). Tractable operations for arithmetic circuits of probabilistic models. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, p. 3943–3951.
- Taghipour, N., Fierens, D., Davis, J., & Blockeel, H. (2013). Lifted variable elimination: Decoupling the operators from the constraint language. *Journal Artificial Intelligence Research*, 47(1), 393–439.
- Team, S. D. (2015). Stan: A c++ library for probability and sampling, version 2.5. 0..
- Van Den Broeck, G. (2011). On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS'11*, p. 1386–1394.
- Van Den Broeck, G., & Darwiche, A. (2013). On the complexity and approximation of binary evidence in lifted inference. In *Proceedings of the 26th International Conference on Neural Information Processing Systems, NIPS'13*, pp. 81–86.
- Van Den Broeck, G., Taghipour, N., Meert, W., Davis, J., & De Raedt, L. (2011). Lifted probabilistic inference by first-order knowledge compilation. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, IJCAI'11*, p. 2178–2185. AAAI Press.
- Vennekens, J., Denecker, M., & Bruynooghe, M. (2009). Cp-logic: A language of causal probabilistic events and its relation to logic programming. *Theory and practice of logic programming*, 9(3), 245–308.
- Vennekens, J., Verbaeten, S., & Bruynooghe, M. (2004). Logic programs with annotated disjunctions. In *20th International Conference on Logic Programming*, Vol. 3132 of *Lecture Notes in Computer Science*, pp. 431–445. Springer.

- Warren, D. S. (1992). Memoing for logic programs. *Communications of the ACM*, 35, 93–111.
- Wielemaker, J., Schrijvers, T., Triska, M., & Lager, T. (2011). Swi-prolog. *Theory and Practice of Logic Programming*, 12, 67 – 96.
- Wu, Y., Srivastava, S., Hay, N., Du, S., & Russell, S. (2018). Discrete-continuous mixtures in probabilistic programming: Generalized semantics and inference algorithms. In *International Conference on Machine Learning*, pp. 5343–5352. PMLR.
- Zeng, Z., & Van Den Broeck, G. (2020). Efficient search-based weighted model integration. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, pp. 175–185. PMLR.
- Zhang, N. L., & Poole, D. L. (1996). Exploiting causal independence in bayesian network inference. *Journal of Artificial Intelligence Research*, 5, 301–328.