

An Extended Semantics for Logic Programs with Annotated Disjunctions and its Efficient Implementation

Fabrizio Riguzzi¹ and Terrance Swift²

¹ ENDIF – University of Ferrara
Via Saragat 1, I-44122, Ferrara, Italy
fabrizio.riguzzi@unife.it

² CENTRIA – Universidade Nova de Lisboa
tswift@cs.suysb.edu

Abstract. Logic Programming with Annotated Disjunctions (LPADs) is a formalism for modeling probabilistic information that has recently received increased attention. The LPAD semantics, while being simple and clear, suffers from the requirement of having function free-programs, which is a strong limitation. In this paper we present an extension of the semantics that removes this restriction and allows us to write programs modeling infinite domains, such as Hidden Markov Models. We show that the semantics is well-defined for a large class of programs. Moreover, we present the algorithm “Probabilistic Inference with Tabling and Answer subsumption” (PITA) for computing the probability of queries to programs according to the extended semantics. Tabling and answer subsumption not only ensure the correctness of the algorithm with respect to the semantics but also make it very efficient on programs without function symbols.

PITA has been implemented in XSB and tested on six domains: two with function symbols and four without. The execution times are compared with those of ProbLog, `cpaint` and CVE. PITA was almost always able to solve larger problems in a shorter time on both type of domains.

1 Introduction

Many real world domains only can be represented effectively if we are able to model uncertainty. Recently, there has been an increased interest in logic languages representing probabilistic information due to their successful use in Machine Learning.

Logic Programs with Annotated Disjunction (LPADs) [21] have attracted the attention of various researchers due to their clarity, simplicity, modeling power and ability to model causation. Their semantics is an instance of the distribution semantics [17]: a theory defines a probability distribution over logic programs and the probability of a query is obtained by summing the probabilities of the programs where the query is true. The semantics of LPADs proposed in [21] requires the programs to be function-free, which is a strong requirement ruling out many interesting programs. Thus, we propose a version of the semantics that allows function symbols, along the lines of [17,12].

The new semantics is based on a program transformation technique that not only allows proving the correctness the semantics but also provides an efficient procedure for computing the probability of queries from LPADs. The algorithm “Probabilistic Inference with Tabling and Answer subsumption” (PITA) builds explanations for every

subgoal encountered during a derivation of the query. The explanations are compactly represented using Binary Decision Diagrams (BDDs) that also allow an efficient computation of the probability. Since all the explanations for a subgoal must be found, tabling is very useful for storing such information. Tabling has already been shown useful for probabilistic logic programming in [6,14,7]. PITA transforms the input LPAD into a normal logic programs in which the subgoals have an extra argument storing a BDD that represents the explanations for its answers. Moreover, we also exploit answer subsumption to combine explanations coming from different clauses.

PITA draws inspiration from [5] that first proposed to use BDDs for computing the probability of queries for the Problog language, a minimalistic probabilistic extension of Prolog, and from [15] that applied BDDs to the more general LPAD syntax. Other approaches for reasoning on LPADs include [14], where SLG resolution is extended by repeatedly branching on disjunctive clauses, and [10], where CVE is presented that transforms an LPAD into an equivalent Bayesian network and then performs inference on the network using the variable elimination algorithm.

PITA was tested on a number of datasets, both with and without function symbols, in order to evaluate its efficiency. The execution times of PITA were compared with those of `cp1int` [15], CVE [10] and ProbLog [8]. PITA was able to successfully solve more complex queries than the other algorithms in most cases and it was also almost always faster both on datasets with and without function symbols.

The paper is organized as follows. Section 2 illustrates the syntax and semantics of LPADs. Section 3 discusses the semantics of LPADs with function symbols. Section 4 gives an introduction to BDDs. Section 5 defines dynamic stratification for LPADs. Section 6 briefly recalls tabling and answer subsumption. Section 7 presents PITA and shows its correctness. Section 8 describes the experiments and Section 9 concludes the paper and presents directions for future works.

2 Logic Programs with Annotated Disjunctions

A *Logic Program with Annotated Disjunctions* [21] consists of a finite set of annotated disjunctive clauses of the form $h_1 : \alpha_1 \vee \dots \vee h_n : \alpha_n \leftarrow b_1, \dots, b_m$. In such a clause h_1, \dots, h_n are logical atoms, b_1, \dots, b_m logical literals, and $\{\alpha_1, \dots, \alpha_n\}$ real numbers in the interval $[0, 1]$ such that $\sum_{j=1}^n \alpha_j \leq 1$. $h_1 : \alpha_1 \vee \dots \vee h_n : \alpha_n$ is called the *head* and b_1, \dots, b_m is called the *body*. Note that if $n = 1$ and $\alpha_1 = 1$ a clause corresponds to a normal program clause, sometimes called a *non-disjunctive* clause. If $\sum_{j=1}^n \alpha_j < 1$, the head of the annotated disjunctive clause implicitly contains an extra atom *null* that does not appear in the body of any clause and whose annotation is $1 - \sum_{j=1}^n \alpha_j$. For a clause C of the form above, we define $head(C)$ as $\{(h_i : \alpha_i) | 1 \leq i \leq n\}$ if $\sum_{i=1}^n \alpha_i = 1$ and as $\{(h_i : \alpha_i) | 1 \leq i \leq n\} \cup \{(null : 1 - \sum_{i=1}^n \alpha_i)\}$ otherwise. Moreover, we define $body(C)$ as $\{b_i | 1 \leq i \leq m\}$, $h_i(C)$ as h_i and $\alpha_i(C)$ as α_i .

If the LPAD is ground, a clause represents a probabilistic choice between the non-disjunctive clauses obtained by selecting only one atom in the head. As usual, if the LPAD T is not ground, T can be assigned a meaning by computing its grounding, $ground(T)$. The semantics of LPADs, given in [21], requires the ground program to be finite, so the program must not contain function symbols if it contains variables.

By choosing a head atom for each ground clause of an LPAD we get a normal logic program called a *possible world* of the LPAD (called an *instance* of the LPAD in [21]). A probability distribution is defined over the space of possible worlds by assuming independence between the choices made for each clause.

More specifically, an *atomic choice* is a triple (C, θ, i) where $C \in T$, θ is a substitution that grounds C and $i \in \{1, \dots, |head(C)|\}$. (C, θ, i) means that, for ground clause $C\theta$, the head $h_i(C)$ was chosen. A set of atomic choices κ is *consistent* if $(C, \theta, i) \in \kappa, (C, \theta, j) \in \kappa \Rightarrow i = j$, i.e., only one head is selected for a ground clause. A *composite choice* κ is a consistent set of atomic choices. The *probability* $P(\kappa)$ of a *composite choice* κ is the product of the probabilities of the individual atomic choices, i.e. $P(\kappa) = \prod_{(C, \theta, i) \in \kappa} \alpha_i(C)$.

A *selection* σ is a composite choice that, for each clause $C\theta$ in $ground(T)$, contains an atomic choice (C, θ, i) in σ . We denote the set of all selections σ of a program T by \mathcal{S}_T . A selection σ identifies a normal logic program w_σ defined as follows $w_\sigma = \{(h_i(C)\theta \leftarrow body(C))\theta \mid (C, \theta, i) \in \sigma\}$. w_σ is called a *possible world* (or simply *world*) of T . \mathcal{W}_T denotes the set of all the possible worlds of T . Since selections are composite choices, we can assign a probability to possible worlds: $P(w_\sigma) = P(\sigma) = \prod_{(C, \theta, i) \in \sigma} \alpha_i(C)$.

We consider only *sound* LPADs, in which every possible world has a total model according to the Well-Founded Semantics (WFS) [20]. In this way, the uncertainty is modeled only by means of the disjunctions in the head and not by the features of the semantics. In the following, $w_\sigma \models \phi$ means that the atom ϕ is true in the well-founded model of the program w_σ .

The probability of a ground atom ϕ according to an LPAD T is given by the sum of the probabilities of the possible worlds where the atom is true under the WFS: $P(\phi) = \sum_{\sigma \in \mathcal{S}_T, w_\sigma \models \phi} P(\sigma)$. It is easy to see that P satisfies the axioms of probability.

Example 1. Consider the dependency of sneezing on having the flu or hay fever:

$$\begin{aligned} C_1 &= strong_sneezing(X) : 0.3 \vee moderate_sneezing(X) : 0.5 \leftarrow flu(X). \\ C_2 &= strong_sneezing(X) : 0.2 \vee moderate_sneezing(X) : 0.6 \leftarrow hay_fever(X). \\ C_3 &= flu(david). \\ C_4 &= hay_fever(david). \end{aligned}$$

This program models the fact that sneezing can be caused by flu or hay fever. The query $strong_sneezing(david)$ is true in 5 of the 9 instances of the program and its probability is

$$P_T(strong_sneezing(david)) = 0.3 \cdot 0.2 + 0.3 \cdot 0.6 + 0.3 \cdot 0.2 + 0.5 \cdot 0.2 + 0.2 \cdot 0.2 = 0.44$$

Even if we assumed independence between the choices for individual ground clauses, this does not represent a restriction, in the sense that this still allow to represent all the joint distributions of atoms of the Herbrand base that are representable with a Bayesian network over those variables. Details of the proof are omitted for lack of space.

LPADs can be written by the user or learned from data. When written by the user, the best approach is to write each clause so that it models a causal mechanism of the domain and to choose the parameters on the basis of his knowledge of the mechanism.

3 A Semantics for LPADs with Function Symbols

If a non-ground LPAD T contains function symbols, then the semantics given in Section 2 is not well-defined. In this case, each possible world w_σ is the result of an infinite number of choices and the probability $P(w_\sigma)$ of w_σ is 0 since it is given by the product of an infinite number of factors all smaller than 1. Thus, the probability of a formula is 0 as well, since it is a sum of terms all equal to 0.

Therefore a new definition of the LPAD semantics is necessary. We provide such a definition following the approach in [12] for assigning a semantics to ICL programs with function symbols. A similar result can be obtained using [17].

A composite choice κ identifies a set of possible worlds ω_κ that contains all the worlds relative to a selection that is a superset of κ , i.e., $\omega_\kappa = \{w_\sigma | \sigma \in \mathcal{S}_T, \sigma \supseteq \kappa\}$. We define the set of possible worlds associated to a set of composite choices K : $\omega_K = \bigcup_{\kappa \in K} \omega_\kappa$.

Given a ground atom ϕ , we define the notion of explanation, covering set of composite choices and mutually incompatible set of explanations. A finite composite choice κ is an *explanation* for ϕ if ϕ is true in every world of ω_κ . In Example 1, the composite choice $\{(C_1, \{X/david\}, 1)\}$ is an explanation for *strong_sneezing(david)*. A set of composite choices K is *covering* with respect to ϕ if every world w_σ in which ϕ is true is such that $w_\sigma \in \omega_K$. In Example 1, the set of composite choices

$$K_1 = \{\{(C_1, \{X/david\}, 1)\}, \{(C_2, \{X/david\}, 1)\}\} \quad (1)$$

is covering for *strong_sneezing(david)*. Two composite choices κ_1 and κ_2 are *incompatible* if their union is inconsistent, i.e., if there exists a clause C and a substitution θ grounding C such that $(C, \theta, j) \in \kappa_1, (C, \theta, k) \in \kappa_2$ and $j \neq k$. A set K of composite choices is *mutually incompatible* if for all $\kappa_1 \in K, \kappa_2 \in K, \kappa_1 \neq \kappa_2 \Rightarrow \kappa_1$ and κ_2 are incompatible. The set of composite choices

$$\begin{aligned} K_2 = \{ & \{(C_1, \{X/david\}, 1), (C_2, \{X/david\}, 2)\}, \\ & \{(C_1, \{X/david\}, 1), (C_2, \{X/david\}, 3)\}, \\ & \{(C_2, \{X/david\}, 1)\} \} \end{aligned} \quad (2)$$

is mutually incompatible for the theory of Example 1. The following results of [12] hold also for LPADs.

- Given a finite set K of finite composite choices, there exists a finite set K' of mutually incompatible finite composite choices such that $\omega_K = \omega_{K'}$.
- If K_1 and K_2 are both mutually incompatible sets of composite choices such that $\omega_{K_1} = \omega_{K_2}$ then $\sum_{\kappa \in K_1} P(\kappa) = \sum_{\kappa \in K_2} P(\kappa)$

Thus, we can define a unique probability measure $\mu : \Omega_T \rightarrow [0, 1]$ where Ω_T is defined as the set of sets of worlds identified by finite sets of finite composite choices: $\Omega_T = \{\omega_K | K \text{ is a finite set of finite composite choices}\}$. It is easy to see that Ω_T is an algebra over \mathcal{W}_T . Then μ is defined by $\mu(\omega_K) = \sum_{\kappa \in K'} P(\kappa)$ where K' is a finite set of finite composite choices that is mutually incompatible and such that $\omega_K = \omega_{K'}$. As for ICL, $\langle \mathcal{W}_T, \Omega_T, \mu \rangle$ is a probability space [9].

Definition 1. The probability of a ground atom ϕ is given by $P(\phi) = \mu(\{w_\sigma | w_\sigma \in \mathcal{W}_T \wedge w_\sigma \models \phi\})$

Theorem 2 in Section 7 shows that, if T is a sound LPAD with bounded term-size and ϕ is a ground atom, there is a finite set K of explanations of ϕ such that K is covering. Therefore $P(\phi)$ is well-defined.

In the case of Example 1, K_2 shown in equation 2 is a covering set of explanations for $sneezing(david, strong)$ that is mutually incompatible, so

$$P(sneezing(david, strong)) = 0.3 \cdot 0.6 + 0.3 \cdot 0.2 + 0.2 = 0.44$$

4 Representing Explanations by Means of Decision Diagrams

In order to represent explanations we can use Multivalued Decision Diagrams [19]. An MDD represents a function $f(\mathbf{X})$ taking Boolean values on a set of multivalued variables \mathbf{X} by means of a rooted graph that has one level for each variable. Each node has one child for each possible value of the multivalued variable associated to the level of the node. The leaves store either 0 or 1. Given values for all the variables \mathbf{X} , an MDD can compute the value of $f(\mathbf{X})$ by traversing the graph starting from the root and returning the value associated to the leaf that is reached.

Given a set of explanations K , we obtain a Boolean function f_K in the following way. Each ground clause $C\theta$ appearing in K is associated to a multivalued variable $X_{C\theta}$ with as many values as atoms in the head of C . Each atomic choice (C, θ, i) is represented by the propositional equation $X_{C\theta} = i$. Equations for a single explanation are conjoined and the conjunctions for the different explanations are disjoined.

The set of explanations in Equation (1) can be represented by the function $f_{K_1}(\mathbf{X}) = (X_{C_1\emptyset} = 1) \vee (X_{C_2\emptyset} = 1)$. An MDD can be obtained from a Boolean function: from f_{K_1} the MDD shown in Figure 1(a) is obtained.




Fig. 1. Decision diagrams for Example 1.

Given a MDD M , we can identify a set of explanations K_M associated to M that is obtained by considering each path from the root to a 1 leaf as an explanation. It is easy to see that K is a set of explanations and M is obtained from f_K , K and K_M represent the same set of worlds, i.e., that $\omega_K = \omega_{K_M}$.

The important role of MDDs is that K_M is mutually incompatible because at each level we branch on a variable and the explanations associated to the leaves that are below a child of a node are incompatible with those of the other children of the node.

By converting a set of explanations into a mutually incompatible set of explanations, MDDs allow to compute $\mu(\omega_K)$ given any K . This is equivalent to computing the probability of a DNF formula which is an NP-hard problem but decision diagrams offer also a practical algorithm that was shown better than other methods [5].

Decision diagrams can be built with various software packages that provide highly efficient implementation of Boolean operations. However, most packages are restricted to work on Binary Decision Diagram (BDD), i.e., decision diagrams where all the variables are Boolean. To work on MDD with a BDD package, we must represent multivalued variables by means of binary variables. Various options are possible, we found that the following, proposed in [4], gives the best performance. For a variable X having n values, we use $n - 1$ Boolean variables X_1, \dots, X_{n-1} and we represent the equation $X = i$ for $i = 1, \dots, n-1$ by means of the conjunction $\overline{X_1} \wedge \overline{X_2} \wedge \dots \wedge \overline{X_{i-1}} \wedge X_i$, and the equation $X = n$ by means of the conjunction $\overline{X_1} \wedge \overline{X_2} \wedge \dots \wedge \overline{X_{n-1}}$. The BDD representation of the function f_{K_1} is given in Figure 1(b). The Boolean variables are associated with the following parameters: $P(X_1) = P(X = 1) \dots P(X_i) = \frac{P(X=i)}{\prod_{j=1}^{i-1} (1-P(X_{j-1}))}$.

5 Dynamic Stratification of LPADs

One of the most important formulations of stratification is that of *dynamic* stratification. [13] shows that a program has a 2-valued well-founded model iff it is dynamically stratified, so that it is the weakest notion of stratification that is consistent with the WFS. As presented in [13], dynamic stratification computes strata via operators on 3-valued interpretations – pairs of the form $\langle T; F \rangle$, where T and F are subsets of the Herbrand base H_P of a normal program P .

Definition 2. For a normal program P , sets T and F of ground atoms, and a 3-valued interpretation I we define

$$\begin{aligned} True_I(T) &= \{A : val_I(A) \neq \tau \text{ and there is a clause } B \leftarrow L_1, \dots, L_n \text{ in } P \text{ and a} \\ &\quad \text{ground substitution } \theta \text{ such that } A = B\theta \text{ and for every } 1 \leq i \leq n \text{ either } L_i\theta \text{ is true} \\ &\quad \text{in } I, \text{ or } L_i\theta \in T\}; \\ False_I(F) &= \{A : val_I(A) \neq \mathbf{f} \text{ and for every clause } B \leftarrow L_1, \dots, L_n \text{ in } P \text{ and} \\ &\quad \text{ground substitution } \theta \text{ such that } A = B\theta \text{ there is some } i (1 \leq i \leq n), \text{ such that } L_i\theta \\ &\quad \text{is false in } I \text{ or } L_i\theta \in F\}. \end{aligned}$$

The conditions $val_I(A) \neq \tau$ and $val_I(A) \neq \mathbf{f}$ are inessential, but ensure that only *new* facts are included in $True_I(T)$ and $False_I(F)$, and simplify the definition of dynamic strata below. [13] shows that $True_I$ and $False_I$ are both monotonic and defines \mathcal{T}_I as the least fixed point of $True_I$ and \mathcal{F}_I as the greatest fixed point of $False_I$. In words, the operator \mathcal{T}_I extends the interpretation I to add the new atomic facts that can be derived from P knowing I ; \mathcal{F}_I adds the new negations of atomic facts that can be shown false in P by knowing I (via the uncovering of unfounded sets). An iterated fixed point operator builds up dynamic strata by constructing successive partial interpretations as follows

Definition 3 (Iterated Fixed Point and Dynamic Strata). For a program P let

$$\begin{aligned} WFM_0 &= \langle \emptyset; \emptyset \rangle; \\ WFM_{\alpha+1} &= WFM_\alpha \cup \langle \mathcal{T}_{WFM_\alpha}; \mathcal{F}_{WFM_\alpha} \rangle; \\ WFM_\alpha &= \bigcup_{\beta < \alpha} WFM_\beta, \text{ for limit ordinal } \alpha. \end{aligned}$$

Let $WFM(P)$ denote the fixed point interpretation WFM_δ , where δ is the smallest countable ordinal such that both sets \mathcal{T}_{WFM_δ} and \mathcal{F}_{WFM_δ} are empty. We refer to δ as the depth of program P . The stratum of atom A , is the least ordinal β such that $A \in WFM_\beta$ (where A may be either in the true or false component of WFM_β).

[13] shows that the iterated fixed point $WFM(P)$ is in fact the well-founded model and that any undefined atoms of the well-founded model do not belong to any stratum – i.e. they are not added to WFM_δ for any ordinal δ .

Dynamic stratification captures the order in which recursive components of a program must be evaluated. Because of this, dynamic stratification is useful for modeling operational aspects of program evaluation. Fixed-order dynamic stratification [16], used in Section 7, replaces the definition of $False_I(F)$ in Definition 2 is by

$$False_I(F) = \{A : val_I(A) \neq \mathbf{f} \text{ and for every clause } B \leftarrow L_1, \dots, L_n \text{ in } P \text{ and ground substitution } \theta \text{ such that } A = B\theta \text{ there exists a **failing prefix**: i.e., there is some } i (1 \leq i \leq n), \text{ such that } L_i\theta \text{ is false in } I \text{ or } L_i\theta \in F, \text{ and for all } j (1 \leq j \leq i - 1), L_j\theta \text{ is true in } I\}.$$

[16] describes how fixed-order dynamic stratification captures those programs that a tabled evaluation can evaluate with a fixed literal selection strategy (i.e. without the SLG operations of SIMPLIFICATION and DELAY). As shown from the following example, fixed-order stratification is a fairly weak condition for a program.

Example 2. The following program has a 2-valued well-founded model and so is dynamically stratified, but does not belong to other stratification classes, such as local, modular, or weak stratification.

$$\begin{array}{ll} s \leftarrow \neg s, p. & s \leftarrow \neg p, \neg q, \neg r. \\ p \leftarrow q, \neg r, \neg s. & q \leftarrow r, \neg p. \\ r \leftarrow p, \neg q. & \end{array}$$

p , q , and r all belong to stratum 0, while s belongs to stratum 1. The simple program

$$p \leftarrow \neg p. \qquad p.$$

is fixed-order stratified, but not locally, modularly, or weakly stratified. Fixed-order stratification is more general than local stratification, and than modular stratification (since modular stratified programs can be decidable rearranged so that they have failing prefixes). It is neither more nor less general than weak stratification.

The above definitions of (fixed-order) dynamic stratification for normal programs can be straightforwardly adapted to LPADs – an LPAD is *(fixed-order) dynamically stratified* if each $w \in \mathcal{W}_T$ is (fixed-order) dynamically stratified.

6 Tabling and Answer Subsumption

The idea behind tabling is to maintain in a table both subgoals encountered in a query evaluation and answers to these subgoals. If a subgoal is encountered more than once, the evaluation reuses information from the table rather than re-performing resolution

against program clauses. Although the idea is simple, it has important consequences. First, tabling ensures termination of programs with the *bounded term-size property*. A program P has the bounded term-size property if there is a finite function $f : N \rightarrow N$ such that if a query term Q to P has size $size(Q)$, then no term used in the derivation of Q has size greater than $f(size(Q))$. This makes it easier to reason about termination than in basic Prolog. Second, tabling can be used to evaluate programs with negation according to the WFS. Third, for queries to wide classes of programs, such as datalog programs with negation, tabling can achieve the optimal complexity for query evaluation. And finally, tabling integrates closely with Prolog, so that Prolog’s familiar programming environment can be used, and no other language is required to build complete systems. As a result, a number of Prologs now support tabling including XSB, YAP, B-Prolog, ALS, and Ciao. In these systems, a predicate p/n is evaluated using SLDNF by default: the predicate is made to use tabling by a declaration such as *table p/n* that is added by the user or compiler.

This paper makes use of a tabling feature called *answer subsumption*. Most formulations of tabling add an answer A to a table for a subgoal S only if A is not a variant (as a term) of any other answer for S . However, in many applications it may be useful to order answers according to a partial order or (upper semi-)lattice. In the case of a lattice, answer subsumption may be specified by means of a declaration such as *table p(.,or/3 - zero/1)*, for an unary predicate p , where *zero/1* is the bottom element of the lattice and *or/3* is the join operation of the lattice. For example, in the PITA algorithm for LPADs presented in Section 7, if a table had an answer $p(a, E_1)$ and a new answer $p(a, E_2)$ were derived, where E_1 and E_2 are probabilistic explanations, the answer $p(a, E_1)$ is replaced by $p(a, E_3)$, where E_3 is obtained by calling $or(E_1, E_2, E_3)$ and is the logical disjunction of the first two explanations, as stored in a BDD³. Answer subsumption over arbitrary upper semi-lattices is implemented in XSB for stratified programs [18]; in addition, the mode-directed tabling of B-Prolog can also be seen as a form of answer subsumption.

Section 7 uses SLG resolution [3] extended with answer subsumption in its proof of Theorem 2, although similar results could be extended to other tabling formalisms that support negation and answer subsumption.

7 Program Transformation

The first step of the PITA algorithm is to apply a program transformation to an LPAD to create a normal normal program that contains calls for manipulating BDDs. In our implementation, these calls provide a Prolog interface to the CUDD⁴ C library and use the following predicates⁵

- *init, end*: for allocation and deallocation of a BDD manager, a data structure used to keep track of the memory for storing BDD nodes;

³ The logical disjunction E_3 can be seen as subsuming E_1 and E_2 over the partial order of implication defined on logical formulas.

⁴ <http://vlsi.colorado.edu/~fabio/>

⁵ BDDs are represented in CUDD as pointers to their root node.

- *zero(-BDD), one(-BDD), and(+BDD1,+BDD2,-BDDO), or(+BDD1,+BDD2,-BDDO), not(+BDD1,-BDDO)*: Boolean operations between BDDs;
- *add_var(+N_Val,+Probs,-Var)*: addition of a new multi-valued variable with *N_Val* values and parameters *Probs*;
- *equality(+Var,+Value,-BDD)*: *BDD* represents *Var=Value*, i.e. that the random variable *Var* is assigned *Value* in the BDD;
- *ret_prob(+BDD,-P)*: returns the probability of the formula encoded by *BDD*.

add_var(+N_Val,+Probs,-Var) adds a new random variable associated to a new instantiation of a rule with *N_Val* head atoms and parameters list *Probs*. The auxiliary predicate *get_var_n(+R,+S,+Probs,-Var)* is used to wrap *add_var/3* and avoid adding a new variable when one already exists for an instantiation. As shown below, a new fact *var(R,S,Var)* is asserted each time a new random variable is created, where *R* is an identifier for the LPAD clause, *S* is a list of constants, one for each variables of the clause, and *Var* is an integer that identifies the random variable associated with clause *R* under a particular grounding. The auxiliary predicates has the following definition

$$\begin{aligned} \text{get_var_n}(R, S, Probs, Var) \leftarrow \\ & (\text{var}(R, S, Var) \rightarrow \text{true}; \\ & \text{length}(Probs, L), \text{add_var}(L, Probs, Var), \text{assert}(\text{var}(R, S, Var))). \end{aligned}$$

where *Probs* is a list of floats that stores the parameters in the head of rule *R*.

The PITA transformation applies to clauses, literals and atoms. If *h* is an atom, *PITAh(h)* is *h* with the variable *BDD* added as the last argument. If *b_j* is an atom, *PITAb(b_j)* is *b_j* with the variable *B_j* added as the last argument. In either case for an atom *a*, *BDD(PITa(a))* is the value of the last argument of *PITa(a)*,

If *b_j* is negative literal $\neg a_j$, *PITAb(b_j)* is the conditional $(PITa'_b(a_j) \rightarrow \text{not}(BN_j, B_j); \text{one}(B_j))$, where *PITa'_b(a_j)* is *a_j* with the variable *BN_j* added as the last argument. In other words the input BDD, *BN_k*, is negated if it exists; otherwise the BDD for the constant function 1 is returned.

A non-disjunctive fact $C_r = h$ is transformed into the clause $PITa(C_r) = PITAh(h) \leftarrow \text{one}(BDD)$.

A disjunctive fact $C_r = h_1 : \alpha_1 \vee \dots \vee h_n : \alpha_n$. where the parameters sum to 1, is transformed into the set of clauses $PITa(C_r)$ ⁶

$$PITa(C_r, 1) = PITAh(h_1) \leftarrow \text{get_var_n}(i, [], [\alpha_1, \dots, \alpha_n], Var), \\ \text{equality}(Var, 1, BDD).$$

$$\dots \\ PITa(C_r, n) = PITAh(h_n) \leftarrow \text{get_var_n}(r, [], [\alpha_1, \dots, \alpha_n], Var), \\ \text{equality}(Var, n, BDD).$$

In the case where the parameters do not sum to one, the clause is first transformed into $h_1 : \alpha_1 \vee \dots \vee h_n : \alpha_n \vee \text{null} : 1 - \sum_1^n \alpha_i$. and then into the clauses above, where the list of parameters is $[\alpha_1, \dots, \alpha_n, 1 - \sum_1^n \alpha_i]$ but the $(n + 1)$ -th clause (the one for *null*) is not generated.

The definite clause $C_r = h \leftarrow b_1, b_2, \dots, b_m$. is transformed into the clause

$$PITa(C_r) = PITAh(h) \leftarrow PITAb(b_1), PITAb(b_2), \text{and}(B_1, B_2, BB_2), \\ \dots, PITAb(b_m), \text{and}(BB_{m-1}, B_m, BDD).$$

⁶ The second argument of *get_var_n* is the empty list because, since we are considering only range restricted programs (cfr. below), a fact does not contain variables.

The disjunctive clause

$$C_r = h_1 : \alpha_1 \vee \dots \vee h_n : \alpha_n \leftarrow b_1, b_2, \dots, b_m.$$

where the parameters sum to 1, is transformed into the set of clauses $PITA(C_r)$

$$PITA(C_r, 1) = PITA_h(h_1) \leftarrow PITA_b(b_1), PITA_b(b_2), \text{and}(B_1, B_2, BB_2), \\ \dots, PITA_b(b_m), \text{and}(BB_{m-1}, B_m, BB_m), \\ \text{get_var_n}(r, VC, [\alpha_1, \dots, \alpha_n], Var), \\ \text{equality}(Var, 1, B), \text{and}(BB_m, B, BDD).$$

$$\dots \\ PITA(C_r, n) = PITA_h(h_n) \leftarrow PITA_b(b_1), PITA_b(b_2), \text{and}(B_1, B_2, BB_2), \\ \dots, PITA_b(b_m), \text{and}(BB_{m-1}, B_m, BB_m), \\ \text{get_var_n}(r, VC, [\alpha_1, \dots, \alpha_n], Var), \\ \text{equality}(Var, n, B), \text{and}(BB_m, B, BDD).$$

where VC is a list containing each variable appearing in C_r . If the parameters do not sum to 1, the same technique used for disjunctive facts is used.

Example 3. Clause C_1 from the LPAD of Example 1 is translated into

$$\text{strong_sneezing}(X, BDD) \leftarrow \text{flu}(X, B_1), \\ \text{get_var_n}(1, [X], [0.3, 0.5, 0.2], Var), \\ \text{equality}(Var, 1, B), \text{and}(B_1, B, BDD). \\ \text{moderate_sneezing}(X, BDD) \leftarrow \text{flu}(X, B_1), \\ \text{get_var_n}(1, [X], [0.3, 0.5, 0.2], Var), \\ \text{equality}(Var, 2, B), \text{and}(B_1, B, BDD).$$

while clause C_3 is translated into

$$\text{flu}(\text{david}, BDD) \leftarrow \text{one}(BDD).$$

In order to answer queries, the goal $\text{solve}(\text{Goal}, P)$ is used, which is defined by

$$\text{solve}(\text{Goal}, P) \leftarrow \text{init}, \text{retractall}(\text{var}(-, -, -)), \\ \text{add_bdd_arg}(\text{Goal}, BDD, \text{GoalBDD}), \\ (\text{call}(\text{GoalBDD}) \rightarrow \text{ret_prob}(BDD, P); P = 0.0), \\ \text{end}.$$

where $\text{add_bdd_arg}(\text{Goal}, BDD, \text{GoalBDD})$ implements $PITA_h(\text{Goal})$. Moreover, various predicates of the LPAD should be declared as tabled. For a predicate p/n , the declaration is `table p(_1, ..., _n, or/3-zero/1)`, that indicates that answer subsumption is used to form the disjunct of multiple explanations: At a minimum, the predicate of the goal should be tabled; as shown in Section 8 it is usually better to table every predicate whose answers have multiple explanations and are going to be reused often.

Correctness of PITA In this section we show two results regarding the PITA transformation and its tabled evaluation⁷. These results ensure on one hand that the semantics is well-defined and on the other hand that the evaluation algorithm is correct. For the purposes of our semantics, we consider the BDDs produced as ground terms, and do not specify them further. We first state the correctness of the PITA transformation with respect to the well-founded semantics of LPADs. Because we allow LPADs to have

⁷ Due to space limitations, our presentation is somewhat informal: a formal presentation with all proofs and supporting definitions can be found at <http://www.ing.unife.it/docenti/FabrizioRiguzzi/Papers/RigSwil10-TR.pdf>.

function symbols, care must be taken to ensure that explanations are finite. To accomplish this, we prove correctness for what we term dynamically-finitary programs, essentially those for which a derivation in the well-founded semantics does not depend on an infinite unfounded set⁸.

Theorem 1 (Correctness of PITA Transformation). Let T be a sound dynamically-finitary LPAD. Then κ is an explanation for a ground atom a iff there is a $PITA_h(a)\theta$ in $WFM(PITA(ground(T)))$, such that κ is a path in $BDD(PITA_h(a)\theta)$ to a 1 leaf.

Theorem 2 below states the correctness of the tabling implementation of PITA, since the BDD returned for a tabled query is the disjunction of a set of covering explanations for that query. The proof uses an extension of SLG evaluation that includes answer subsumption but that is restricted to fixed-order dynamically stratified programs [18], a formalism that models the implementation tested in Section 8. Note that unlike Theorem 1, Theorem 2 does not require the program T to be grounded. However, Theorem 2 does require T to be range restricted in order to ensure that tabled evaluation grounds answers. A normal program/LPAD is *range restricted* if all the variables appearing in the head of each clause appear also in the body. If a normal program is range restricted, every successful SLDNF-derivation for G completely grounds G [11], a result that can be straightforwardly extended to tabled evaluations. In addition, Theorem 2 requires T to have the bounded term-size property (cf. Section 6) to ensure termination and finite explanations.

Theorem 2 (Correctness of PITA Evaluation). Let T be a range restricted, bounded term-size, fixed-order dynamically stratified LPAD and a a ground atom. Let \mathcal{E} be an SLG evaluation of $PITA_h(a)$ against $PITA(T)$, such that answer subsumption is declared on $PITA_h(a)$ using BDD-disjunction. Then \mathcal{E} terminates with an answer ans for $PITA_h(a)$ and $BDD(ans)$ represents a covering set of explanations for a .

Thus range restricted, bounded term-size and fixed-order dynamically stratified LPADs have a finite set of explanations that are covering for a ground atom, so the semantics with function symbols is well-defined.

8 Experiments

PITA was tested on two datasets that contain function symbols: the first is taken from [21] and encodes a Hidden Markov Model (HMM) while the latter from [5] encodes biological networks. Moreover, it was also tested on the four testbeds of [10] that do not contain function symbols. PITA was compared with the exact version of ProbLog⁹ [5] available in the git version of Yap as of 19/12/2009, with the version of `cp1int`

⁸ Dynamically-finitary programs are a strict superclass of the finitary programs of [1] and are neither a subclass nor a superclass of the finitely ground programs of [2]. The formal definition of dynamically-finitary programs is in the full version of this paper.

⁹ ProbLog was not tested on programs with more than two atoms in the head because the publicly available version is not yet able to deal with non-binary variables.

[15] available in Yap 6.0 and with the version of CVE [10] available in ACE-ilProlog 1.2.20¹⁰.

The first problem models a HMM with three states 1, 2 and 3 of which 3 is an end state. This problem is encoded by the program

$$s(0,1):1/3 \vee s(0,2):1/3 \vee s(0,3):1/3.$$

$$s(T,1):1/3 \vee s(T,2):1/3 \vee s(T,3):1/3 \leftarrow T1 \text{ is } T-1, T1 \geq 0, s(T1,F), \setminus + s(T1,3)$$

For this experiment, we query the probability of the HMM being in state 1 at time N for increasing values of N , i.e., we query the probability of $s(N,1)$. In PITA, we did not use reordering of BDDs variables¹¹. The execution times of PITA, CVE and `cpLint` are shown in Figure 2. In this problem tabling provides an impressive speedup, since computations can be reused often.




Fig. 2. Three sided die.

The biological network problems compute the probability of a path in a large graph in which the nodes encode biological entities and the links represents conceptual relations among them. Each programs in this dataset contains a definition of path plus a number of links represented by probabilistic facts. The programs have been sampled from a very large graph and contain 200, 400, . . . , 5000 edges. Sampling was repeated ten times, to obtain a series of 10 programs of increasing size. In each test we queried the probability that the two genes HGNC_620 and HGNC_983 are related. We used the definition of path of [8] that performs loop checking explicitly by keeping the list of visited nodes:

$$\begin{aligned} path(X, Y) &\leftarrow path(X, Y, [X], Z). \\ path(X, Y, V, [Y|V]) &\leftarrow edge(X, Y). \\ path(X, Y, V0, V1) &\leftarrow edge(X, Z), append(V0, _S, V1), \\ &\quad \setminus + member(Z, V0), path(Z, Y, [Z|V0], V1). \end{aligned}$$

We used this definition because it gave better results than the one without explicit loop checking. The possibility of using lists (that require function symbols) allowed in this case more modeling freedom. The predicates $path/2$ and $edge/2$ are tabled.

We ran PITA, ProbLog and `cpLint` on the graphs in sequence starting from the smallest program and in each case we stopped after one day or at the first graph for

¹⁰ All experiments were performed on Linux machines with an Intel Core 2 Duo E6550 (2333 MHz) processor and 4 GB of RAM.

¹¹ For each experiment, we used either group sift automatic reordering or no reordering of BDDs variables depending on which gave the best results.




Fig. 3. Biological graph experiments.

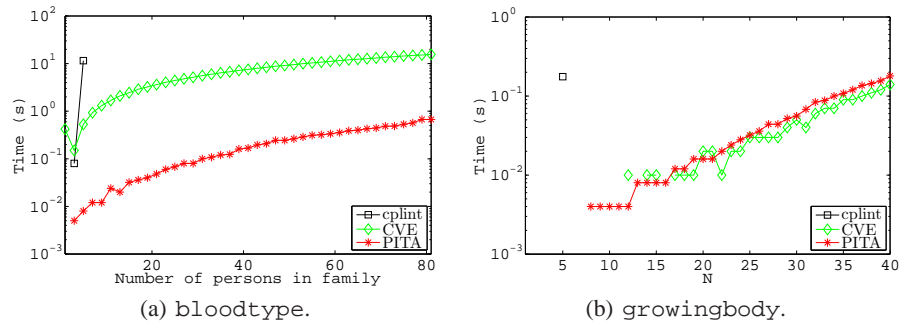


Fig. 4. Datasets from (Meert et al. 2009).

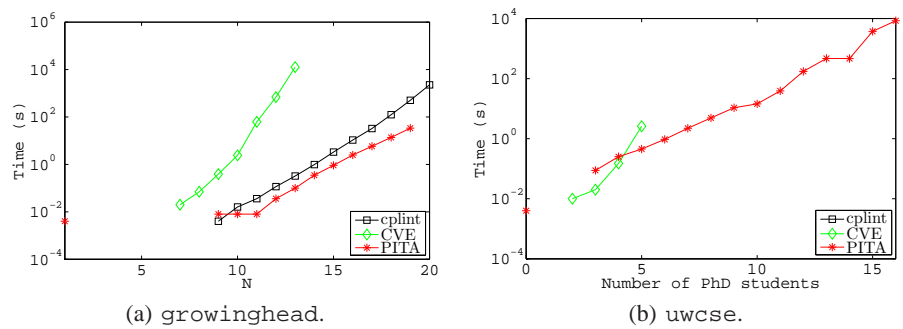


Fig. 5. Datasets from (Meert et al. 2009).

which the program ended for lack of memory¹². In PITA, we used group sift reordering of BDDs variables. Figure 3(a) shows the number of subgraphs for which each algorithm was able to answer the query as a function of the size of the subgraphs, while Figure 3(b) shows the execution time averaged over all and only the subgraphs for which all the algorithms succeeded. PITA was able to solve more subgraphs and in a shorter time than `cplint` and `ProbLog`. For PITA the vast majority of time for larger graphs was spent on BDD maintenance. `ProbLog` ended for lack of memory in three cases out of ten, PITA in two and `cplint` in four. This shows that, even if tabling consumes more memory when finding the explanations, BDDs are built faster and using less memory, probably due to the fact that tabling allows less redundancy (only one BDD is stored for an answer) and a bottom-up construction of the BDDs, which is usually better. This shows that one should table every predicate whose answer have multiple explanations, as *path/2* and *edge/2* above.

The four datasets of [10], served as a final suite of benchmarks. `bloodtype` encodes the genetic inheritance of blood type, `growingbody` contains programs with growing bodies, `growinghead` contains programs with growing heads and `uwcase` encodes a university domain. In PITA we disabled automatic reordering of BDDs variables for all datasets except for `uwcase`. The execution times of `cplint`, `CVE` and PITA are shown in Figures 4(a) and 4(b), 5(a) and 5(b)¹³. PITA was faster than `cplint` in all domains and faster than `CVE` in all domains except `growingbody`. `growingbody`, however, is a domain in which all the clauses are mutually exclusive, thus making possible to compute the probability even without BDDs.

9 Conclusion and Future Works

This paper has made two contributions. The first, semantic, contribution extends LPADs to include functions. By way of proving correctness of the PITA transformation we also characterize those extended LPAD programs whose derived atoms have only finite explanations (dynamically-finitary LPADs); by way of proving correctness of PITA evaluation we characterize those that have only finite sets of explanations (LPADs with the bounded term-size property). Such results ensure that the semantics with function symbols is well-defined.

The PITA transformation also provides a practical reasoning algorithm that was directly used in the experiments of Section 8. The experiments substantiate the PITA approach. Accordingly PITA programs should be easily portable to other tabling engines such as that of `YAP`, `Ciao` and `B Prolog` if they support answer subsumption over general semi-lattices.

In the future, we plan to extend PITA to the whole class of sound LPADs by implementing the `SLG DELAYING` and `SIMPLIFICATION` operations for answer subsumption. In addition, we are developing a version of PITA that is able to answer queries in an approximate way, similarly to [8].

¹² `CVE` was not applied to this dataset because the current version can not handle graph cycles.

¹³ For the missing points at the beginning of the lines a time smaller than 10^{-6} was recorded. For the missing points at the end of the lines the algorithm exhausted the available memory.

References

1. Baselice, S., Bonatti, P., Criscuolo, G.: On finitely recursive programs. *Theory and Practice of Logic Programming* 9(2), 213–238 (2009)
2. Calimeri, F., Cozza, S., Ianni, G., Leone, N.: Computable functions in ASP: Theory and implementation. In: *ICLP*. pp. 407–424 (2008)
3. Chen, W., Warren, D.S.: Tabled evaluation with delaying for general logic programs. *J. ACM* 43(1), 20–74 (1996)
4. De Raedt, L., Demoen, B., Fierens, D., Gutmann, B., Janssens, G., Kimmig, A., Landwehr, N., Mantadelis, T., Meert, W., Rocha, R., Santos Costa, V., Thon, I., Vennekens, J.: Towards digesting the alphabet-soup of statistical relational learning. In: *NIPS*2008 Workshop on Probabilistic Programming* (2008)
5. De Raedt, L., Kimmig, A., Toivonen, H.: ProbLog: A probabilistic Prolog and its application in link discovery. In: *International Joint Conference on Artificial Intelligence*. pp. 2462–2467 (2007)
6. Kameya, Y., Sato, T.: Efficient EM learning with tabulation for parameterized logic programs. In: *Computational Logic*. LNCS, vol. 1861, pp. 269–284. Springer (2000)
7. Kimmig, A., Gutmann, B., Santos Costa, V.: Trading memory for answers: Towards tabling ProbLog. In: *International Workshop on Statistical Relational Learning*. KU Leuven (2009)
8. Kimmig, A., Santos Costa, V., Rocha, R., Demoen, B., De Raedt, L.: On the efficient execution of ProbLog programs. In: *ICLP*. LNCS, vol. 5366, pp. 175–189. Springer (2008)
9. Kolmogorov, A.N.: *Foundations of the Theory of Probability*. Chelsea Publishing Company, New York (1950)
10. Meert, W., Struyf, J., Blockeel, H.: CP-Logic theory inference with contextual variable elimination and comparison to BDD based inference methods. In: *International Conference on Inductive Logic Programming*. KU Leuven, Leuven, Belgium (2009)
11. Muggleton, S.: Learning stochastic logic programs. *Electron. Trans. Artif. Intell.* 4(B), 141–153 (2000)
12. Poole, D.: Abducing through negation as failure: stable models within the independent choice logic. *J. Log. Program.* 44(1-3), 5–35 (2000)
13. Przymusiński, T.: Every logic program has a natural stratification and an iterated least fixed point model. In: *Symposium on Principles of Database Systems*. pp. 11–21. ACM Press (1989)
14. Riguzzi, F.: Inference with logic programs with annotated disjunctions under the well founded semantics. In: *ICLP*. pp. 667–771. No. 5366 in LNCS, Springer (2008)
15. Riguzzi, F.: A top down interpreter for LPAD and CP-logic. In: *Congress of the Italian Association for Artificial Intelligence*. pp. 109–120. No. 4733 in LNAI, Springer (2007)
16. Sagonas, K., Swift, T., Warren, D.S.: The limits of fixed-order computation. *Theor. Comput. Sci.* 254(1-2), 465–499 (2000)
17. Sato, T.: A statistical learning method for logic programs with distribution semantics. In: *ICLP*. pp. 715–729 (1995)
18. Swift, T.: Tabling for non-monotonic programming. *Ann. Math. Artif. Intell.* 25(3-4), 201–240 (1999)
19. Thayse, A., Davio, M., Deschamps, J.P.: Optimization of multivalued decision algorithms. In: *International Symposium on Multiple-Valued Logic*. pp. 171–178 (1978)
20. van Gelder, A., Ross, K., Schlipf, J.: Unfounded sets and well-founded semantics for general logic programs. *J. ACM* 38(3), 620–650 (1991)
21. Vennekens, J., Verbaeten, S., Bruynooghe, M.: Logic programs with annotated disjunctions. In: *ICLP*. LNCS, vol. 3131, pp. 195–209. Springer (2004)