

# COALGEBRAIC SEMANTICS FOR PROBABILISTIC LOGIC PROGRAMMING

TAO GU AND FABIO ZANASI

University College London  
*e-mail address:* tao.gu.18@ucl.ac.uk  
*e-mail address:* f.zanasi@ucl.ac.uk

**ABSTRACT.** Probabilistic logic programming is increasingly important in artificial intelligence and related fields as a formalism to reason about uncertainty. It generalises logic programming with the possibility of annotating clauses with probabilities. This paper proposes a coalgebraic semantics on probabilistic logic programming. Programs are modelled as coalgebras for a certain functor  $F$ , and two semantics are given in terms of cofree coalgebras. First, the cofree  $F$ -coalgebra yields a semantics in terms of derivation trees. Second, by embedding  $F$  into another type  $G$ , as cofree  $G$ -coalgebra we obtain a ‘possible worlds’ interpretation of programs, from which one may recover the usual distribution semantics of probabilistic logic programming. Furthermore, we show that a similar approach can be used to provide a coalgebraic semantics to weighted logic programming.

## 1. INTRODUCTION

Probabilistic logic programming (PLP) [NS92, Dan92, Sat95] is a family of approaches extending the declarative paradigm of logic programming with the possibility of reasoning about uncertainty. This has been proven useful in various applications, including bioinformatics [DRKT07a, MH12], robotics [TBF05] and the semantic web [Zes17].

The most common version of PLP — on which for instance **ProbLog** is based [DRKT07a], the probabilistic analogue of **Prolog** — is defined by annotating clauses in programs with mutually independent probabilities. As for the interpretation, *distribution semantics* [Sat95] is typically used as a benchmark for the various implementations of PLP, such as **pD**, **PRISM** and **ProbLog** [RS14]. While in logic programming the central task is whether a goal is provable using the clauses in the program as axioms, for PLP the typical question one asks is what is the probability of a goal being provable. In distribution semantics, all the clauses are seen as independent random events, and such probability is obtained as the sum of the probabilities of all the *possible worlds* (sets of clauses) in which the goal is provable. The distribution semantics is particularly interesting because it is compatible with the encoding of Bayesian networks as probabilistic logic programs [RS14], thus indicating that PLP can be effectively employed for Bayesian reasoning.

*Key words and phrases:* probabilistic logic programming, distribution semantics, weighted logic programming.

The main goal of this work is to present a coalgebraic perspective on PLP and its distribution semantics. We first consider the case of ground programs (Section 3), that is, programs without variables. Our approach is based on the observation — inspired by the coalgebraic treatment of ‘pure’ logic programming [KMP10] — that ground programs are in 1-1 correspondence with coalgebras for the functor  $\mathcal{M}_{pr}\mathcal{P}_f$ , where  $\mathcal{M}_{pr}$  is the finite multiset functor on  $[0, 1]$  and  $\mathcal{P}_f$  is the finite powerset functor. We then provide two coalgebraic semantics for ground PLP.

- The first interpretation  $\llbracket - \rrbracket$  is in terms of execution trees called *stochastic derivation trees*, which represent parallel SLD-derivations of a program on a goal. Stochastic derivation trees are the elements of the cofree  $\mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra on a given set of atoms  $\text{At}$ , meaning that any goal  $A \in \text{At}$  can be given a semantics in terms of the corresponding stochastic derivation tree by the universal map  $\llbracket - \rrbracket$  to the cofree coalgebra.
- The second interpretation  $\langle\langle - \rangle\rangle$  recovers the usual distribution semantics of PLP. This requires some work, as expressing probability distributions on the possible worlds needs a different coalgebra type. We introduce *distribution trees*, a tree-like representation of the distribution semantics, as the elements of the cofree  $\mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$ -coalgebra on  $\text{At}$ , where  $\mathcal{D}_{\leq 1}$  is the sub-probability distribution monad. In order to characterise  $\langle\langle - \rangle\rangle$  as the map given by universal property of distribution trees, we need a canonical extension of PLP to the setting of  $\mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$ -coalgebras. This is achieved via a ‘possible worlds’ natural transformation  $\mathcal{M}_{pr}\mathcal{P}_f \Rightarrow \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$ .

In the second part of the paper (Section 4) we recover the same framework for arbitrary probabilistic logic programs, possibly including variables. The encoding of programs as coalgebras is subtler. The space of atoms is now a presheaf indexed by a ‘Lawvere theory’ of terms and substitutions. The coalgebra map can be defined in different ways, depending on the substitution mechanism on which one wants to base resolution. For pure logic programs, the definition by term-matching is the best studied, with [KP11] observing that moving from sets to posets is required in order for the corresponding coalgebra map to be well-defined as a natural transformation between presheaves. A different route is taken in [BZ15], where the problem of naturality is neutralised via ‘saturation’, a categorical construction which amounts to defining resolution by unification instead of term-matching.

In developing a coalgebraic treatment of PLP with variables, we follow the saturation route, as it also allows to recover the term-matching approach, via ‘desaturation’ [BZ15]. This provides a cofree coalgebra semantics  $\llbracket - \rrbracket$  for arbitrary PLP programs, as a rather straightforward generalisation of the saturated semantics for pure logic programs. On the other hand, extending the ground distribution semantics  $\langle\langle - \rangle\rangle$  to arbitrary PLP programs poses some challenges: we need to ensure that, in computing the distribution over possible worlds associated to each sub-goal in the computation, each clause of the program is ‘counted’ only once. This is solved by tweaking the coalgebra type of the distribution trees for arbitrary PLP programs, so that some nodes are labelled with clauses of the program. Thanks to this additional information, the term-matching distribution semantics of an arbitrary PLP goal is computable from its distribution tree.

After developing our framework for PLP, in the last part of the paper (Section 5) we show how the same approach yields a coalgebraic semantics for weighted logic programming (WLP). WLP generalises standard logic programming by adding weights to clauses; it is mostly used in the specification of dynamic programming algorithms in various fields, including natural language processing [EB20] and computational biology [DEKM98]. Our coalgebraic

description yields a derivation semantics  $\llbracket - \rrbracket$  both for ground and arbitrary WLP programs, from which one may compute the weight associated with a goal.

In light of the coalgebraic treatment of pure logic programming [KMP10, KP11, BZ13, BZ15], the generalisation to PLP and WLP may not appear so surprising. In fact, we believe its importance is two-fold. First, whereas the derivation semantics  $\llbracket - \rrbracket$  is a straight generalisation of the pure setting, the distribution semantics  $\langle\langle - \rangle\rangle$  is genuinely novel, and does not have counterparts in pure logic programming. Second, this work provides a starting point for a generic coalgebraic treatment of variations of logic programming:

- The coalgebraic approach to pure logic programming has been used as a formal justification [KPS16, KL17] for coinductive logic programming [KL18, GBM<sup>+</sup>07]. Coinduction in the context of probabilistic and weighted logic programs is, to the best of our knowledge, a completely unexplored field, for which the current paper establishes semantic foundations.
- As mentioned, reasoning in Bayesian networks can be seen as a particular case of PLP, equipped with the distribution semantics. Our coalgebraic perspective thus readily applies to Bayesian reasoning, paving the way for combination with recent works [JKZ19, JZ19, DDGK16] modelling belief revision, causal inference and other Bayesian tasks in algebraic terms.

We leave the exploration of these venues as follow-up work.

This work extends the conference paper [GZ19] with the addition of the missing proofs (Appendix D), and novel material on weighted logic programming (Section 5 and Appendix C).

## 2. PRELIMINARIES

**Signature, Terms, and Categories.** A *signature*  $\Sigma$  is a set of function symbols, each equipped with a fixed finite arity. Throughout this paper we fix a signature  $\Sigma$ , and a countably infinite set of variables  $Var = \{x_1, x_2, \dots\}$ . The  $\Sigma$ -terms over  $Var$  are defined as usual. A *context* is a finite sequence of variables  $\langle x_1, x_2, \dots, x_n \rangle$ . With some abuse of notation, we shall often use  $n$  to denote this context. We say a  $\Sigma$ -term  $t$  is *compatible* with context  $n$  if the variables appearing in  $t$  are all contained in  $\{x_1, \dots, x_n\}$ .

We are going to reason about  $\Sigma$ -terms categorically using Lawvere theories. First, we will use  $\mathbf{Ob}(\mathbf{C})$  to denote the set of objects and  $\mathbf{C}[C, D]$  for the set of morphisms  $C \rightarrow D$  in a category  $\mathbf{C}$ . A  $\mathbf{C}$ -indexed *presheaf* is a functor  $F: \mathbf{C} \rightarrow \mathbf{Sets}$ .  $\mathbf{C}$ -indexed presheaves and natural transformations between them form a category  $\mathbf{Sets}^{\mathbf{C}}$ . Recall that the (opposite) Lawvere Theory of  $\Sigma$  is the category  $\mathbf{L}_{\Sigma}^{\text{op}}$  with objects the natural numbers and morphisms  $\mathbf{L}_{\Sigma}^{\text{op}}[n, m]$  the  $n$ -tuples  $\langle t_1, \dots, t_n \rangle$ , where each  $t_i$  is a  $\Sigma$ -term in context  $m$ . For modelling logic programming, it is convenient to think of each  $n \in \mathbf{Ob}(\mathbf{L}_{\Sigma}^{\text{op}})$  as representing the context  $\langle x_1, \dots, x_n \rangle$ , and a morphism  $\langle t_1, \dots, t_n \rangle: n \rightarrow m$  as the substitution transforming  $\Sigma$ -terms in context  $n$  to  $\Sigma$ -terms in context  $m$  by replacing each  $x_i$  with  $t_i$ . For this reason we shall also refer to  $\mathbf{L}_{\Sigma}^{\text{op}}$  morphisms simply as substitutions (notation  $\theta, \tau, \sigma, \dots$ ).

**Logic programming.** We now recall the basics of (pure) logic programming, and refer the reader to [Llo87] for a more systematic exposition. An *alphabet*  $\mathcal{A}$  consists of a signature  $\Sigma$ , a set of variables  $Var$ , and a set of predicate symbols  $\{P_1, P_2, \dots\}$ , each with a fixed finite arity. Given an  $n$ -ary predicate symbol  $P$  in  $\mathcal{A}$ , and  $\Sigma$ -terms  $t_1, \dots, t_n$ ,  $P(t_1 \cdots t_n)$  is called an *atom* over  $\mathcal{A}$ . We use  $A, B, \dots$  to denote atoms. Given an atom  $A$  in context  $n$ , and a substitution  $\theta = \langle t_1, \dots, t_n \rangle: n \rightarrow m$ , we write  $A\theta$  for the *substitution instance* of  $A$  obtained by replacing each appearance of  $x_i$  with  $t_i$  in  $A$ . For convenience, we also use

$\{B_1, \dots, B_k\}\theta$  as a shorthand for  $\{B_1\theta, \dots, B_k\theta\}$ . Given two atoms  $A$  and  $B$  (over  $\mathcal{A}$ ), a *unifier* of  $A$  and  $B$  is a pair  $\langle \sigma, \tau \rangle$  of substitutions such that  $A\sigma = B\tau$ . *Term matching* is a special case of unification, where  $\sigma$  is the identity substitution. In this case we say that  $\tau$  matches  $B$  with  $A$  if  $A = B\tau$ .

A (pure) logic program  $\mathbb{L}$  consists of a finite set of clauses  $\mathcal{C}$  in the form  $H \leftarrow B_1, \dots, B_k$ , where  $H, B_1, \dots, B_k$  are atoms.  $H$  is called the *head* of  $\mathcal{C}$ , and  $B_1, \dots, B_k$  form the *body* of  $\mathcal{C}$ . We denote  $H$  by  $\text{Head}(\mathcal{C})$ , and  $\{B_1, \dots, B_k\}$  by  $\text{Body}(\mathcal{C})$ . A clause  $\mathcal{C}$  with empty body is also called a *fact*. A *goal* is simply an atom, and we use this terminology in the context of certain logic programming reasoning tasks. Since one can regard a clause  $H \leftarrow B_1, \dots, B_k$  as the logic formula  $B_1 \wedge \dots \wedge B_k \rightarrow H$ , we say that a goal  $G$  is *derivable* in  $\mathbb{L}$  if there exists a derivation of  $G$  with empty assumption using the clauses in  $\mathbb{L}$ .

The central task of logic programming is to check whether an atom  $G$  is *provable* in a program  $\mathbb{L}$ , in the sense that some substitution instance of  $G$  is derivable in  $\mathbb{L}$ . The key algorithm for this task is SLD-resolution, see e.g. [Llo87]. We use the notation  $\mathbb{L} \vdash G$  to mean that  $G$  is provable in  $\mathbb{L}$ .

**Probabilistic logic programming.** We now recall the basics of PLP; the reader may consult [DRKT07b, DRKT07a] for a more comprehensive introduction. A probabilistic logic program  $\mathbb{P}$  based on a logic program  $\mathbb{L}$  assigns a probability label  $r$  to each clause  $\mathcal{C}$  in  $\mathbb{L}$ , denoted as  $\text{Label}(\mathcal{C})$ . One may also regard  $\mathbb{P}$  as a set of probabilistic clauses of the form  $r :: \mathcal{C}$ , where  $\mathcal{C}$  is a clause in  $\mathbb{L}$ , and each clause  $\mathcal{C}$  is assigned a unique probability label  $r$  in  $\mathbb{P}$ . We also refer to  $r :: \mathcal{C}$  simply as clauses.

**Example 2.1.** As our leading example we introduce the following probabilistic logic program  $\mathbb{P}^{al}$ . It models the scenario of Mary's house alarm, which is supposed to detect burglars, but it may be accidentally triggered by an earthquake. Mary may hear the alarm if she is awake, but even if the alarm is not sounding, in case she experiences an auditory hallucination (paracusia). The language of  $\mathbb{P}^{al}$  includes 0-ary predicates `Alarm`, `Eearthquake`, `Burglary`, and `Paracusia`, and 1-ary predicates `Wake(-)`, `Hear_alarm(-)` and `Paracusia(-)`, and signature  $\Sigma_{al} = \{\text{Mary}^0\}$  consisting of a constant. We do not have variables here, so  $\mathbb{P}^{al}$  is a ground program. For readability we abbreviate `Mary` as `M` in the program.

0.01 ::	<code>Earthquake</code>	$\leftarrow$	0.01 ::	<code>Paracusia(M)</code>	$\leftarrow$
0.2 ::	<code>Burglary</code>	$\leftarrow$	0.6 ::	<code>Wake(M)</code>	$\leftarrow$
0.5 ::	<code>Alarm</code>	$\leftarrow$ <code>Earthquake</code>	0.8 ::	<code>Hear_alarm(M)</code>	$\leftarrow$ <code>Alarm</code> , <code>Wake(M)</code>
0.9 ::	<code>Alarm</code>	$\leftarrow$ <code>Burglary</code>	0.3 ::	<code>Hear_alarm(M)</code>	$\leftarrow$ <code>Paracusia(M)</code>

As a generalisation of the pure case, in probabilistic logic programming one is interested in the *probability* of a goal  $G$  being provable in a program  $\mathbb{P}$ . There are potentially multiple ways to define such probability — in this paper we focus on Sato's *distribution semantics* [Sat95] as below.

Given a probabilistic logic program  $\mathbb{P} = \{p_1 :: \mathcal{C}_1, \dots, p_n :: \mathcal{C}_n\}$ , let  $|\mathbb{P}|$  be its underlying pure logic program, namely  $|\mathbb{P}| = \{\mathcal{C}_1, \dots, \mathcal{C}_n\}$ . A *sub-program*  $\mathbb{L}$  of  $|\mathbb{P}|$  is a logic program consisting of a subset of the clauses in  $|\mathbb{P}|$ . This justifies using  $\mathcal{P}(|\mathbb{P}|)$  to denote the set of all sub-programs of  $|\mathbb{P}|$ , and using  $\mathbb{L} \subseteq |\mathbb{P}|$  to denote that  $\mathbb{L}$  is a sub-program of  $\mathbb{P}$ . The central concept of the distribution semantics is that  $\mathbb{P}$  determines a distribution  $\mu_{\mathbb{P}}$  over the

sub-programs  $\mathcal{P}(|\mathbb{P}|)$ : for any  $\mathbb{L} \in \mathcal{P}(|\mathbb{P}|)$ ,

$$\mu_{\mathbb{P}}(\mathbb{L}) := \left( \prod_{\mathcal{C}_i \in \mathbb{L}} p_i \right) \cdot \left( \prod_{\mathcal{C}_j \in |\mathbb{P}| \setminus \mathbb{L}} (1 - p_j) \right)$$

where each  $p_i$  is the probability label of the clause  $\mathcal{C}_i$  in  $\mathbb{P}$ . We simply refer to the value  $\mu_{\mathbb{P}}(\mathbb{L})$  as the *probability* of the sub-program  $\mathbb{L}$ . For an arbitrary goal  $G \in \text{At}$ , the *success probability* (or simply the probability)  $\text{Pr}_{\mathbb{P}}(G)$  of  $G$  w.r.t. program  $\mathbb{P}$  is then defined as the sum of the probabilities of all the sub-programs of  $\mathbb{P}$  in which  $G$  is provable:

$$\text{Pr}_{\mathbb{P}}(G) := \sum_{|\mathbb{P}| \supseteq \mathbb{L} \vdash G} \mu_{\mathbb{P}}(\mathbb{L}) = \sum_{|\mathbb{P}| \supseteq \mathbb{L} \vdash G} \left( \prod_{\mathcal{C}_i \in \mathbb{L}} p_i \cdot \prod_{\mathcal{C}_j \in |\mathbb{P}| \setminus \mathbb{L}} (1 - p_j) \right) \quad (2.1)$$

Intuitively one can regard every clause in  $\mathbb{P}$  as a random event, then every sub-program  $\mathbb{L}$  can be seen as a possible world over these events, and  $\mu_{\mathbb{P}}$  is a distribution over all the possible worlds. The success probability  $\text{Pr}_{\mathbb{P}}(G)$  is then simply the probability of all the possible worlds in which the goal  $G$  is provable.

**Example 2.2.** For the program  $\mathbb{P}^{al}$ , consider the goal `Hear_alarm(M)`. By the definition of distribution semantics (2.1), we can compute that the success probability  $\text{Pr}_{\mathbb{P}^{al}}(\text{Hear\_alarm}(\text{M}))$  is  $\approx 0.0911$ .

### 3. GROUND CASE

In this section we introduce a coalgebraic semantics for *ground* probabilistic logic programming, i.e. for those programs where no variable appears. Our approach consists of two parts. First, we represent PLP programs as coalgebras (Subsection 3.1) and their executions as a final coalgebra semantics (Subsection 3.2) — this is a straightforward generalisation of the coalgebraic treatment of pure logic programs given in [KMP10]. Next, in Subsection 3.3 we investigate how to represent the distribution semantics as a final coalgebra, via a transformation of the coalgebra type of PLP programs. Appendix A shows how the success probability of a goal is effectively computable from the above representation.

**3.1. Coalgebraic Representation of PLP.** A ground PLP program can be represented as a coalgebra for the composite  $\mathcal{M}_{pr}\mathcal{P}_f: \mathbf{Sets} \rightarrow \mathbf{Sets}$  of the finite probability functor  $\mathcal{M}_{pr}: \mathbf{Sets} \rightarrow \mathbf{Sets}$  and the finite powerset functor  $\mathcal{P}_f: \mathbf{Sets} \rightarrow \mathbf{Sets}$ . The definition of  $\mathcal{M}_{pr}$  deserves some further explanation. It can be seen as the finite multiset functor based on the commutative monoid  $([0, 1], 0, \vee)$ , where  $\vee$  is the probabilistic ‘or’ for independent events defined as  $a \vee b := 1 - (1 - a)(1 - b)$ . That is to say, on objects,  $\mathcal{M}_{pr}(A)$  is the set of all *finite probability assignments*  $\varphi: A \rightarrow [0, 1]$ , namely those  $\varphi$  with a finite support  $\text{supp}(\varphi) := \{a \in A \mid \varphi(a) \neq 0\}$ . For  $\varphi$  with support  $\{a_1, \dots, a_k\}$  and values  $\varphi(a_i) = r_i$ , it will often be convenient to use the standard notation  $\varphi = r_1 a_1 + \dots + r_k a_k$  or  $\varphi = \sum_{i=1}^k r_i a_i$ , where the purely formal ‘+’ and ‘ $\sum$ ’ here should not be confused with the arithmetic addition. On morphisms,  $\mathcal{M}_{pr}(h: A \rightarrow B)$  maps  $\sum_{i=1}^k r_i a_i$  to  $\sum_{i=1}^k r_i h(a_i)$ .

Fix a ground probabilistic logic program  $\mathbb{P}$  on a set of ground atoms  $\text{At}$ . The definition of  $\mathbb{P}$  can be encoded as an  $\mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra  $p: \text{At} \rightarrow \mathcal{M}_{pr}\mathcal{P}_f(\text{At})$ , as follows. Given  $A \in \text{At}$ ,

$$p(A): \quad \mathcal{P}_f(\text{At}) \quad \rightarrow \quad [0, 1]$$

$$\{B_1, \dots, B_n\} \quad \mapsto \quad \begin{cases} r & \text{if } r :: A \leftarrow B_1, \dots, B_n \text{ is a clause in } \mathbb{P} \\ 0 & \text{otherwise.} \end{cases}$$

Or, equivalently,  $p(A) := \sum_{(r :: A \leftarrow B_1, \dots, B_n) \in \mathbb{P}} r\{B_1, \dots, B_n\}$ . Note that each  $p(A)$  has a finite support because the program  $\mathbb{P}$  consists of finitely many clauses.

**Example 3.1.** Consider program  $\mathbb{P}^{al}$  from Example 2.1. The set of ground atoms  $\text{At}_{al}$  is  $\{\text{Alarm}, \text{Earthquake}, \text{Burgary}, \text{Wake}(\text{M}), \text{Paracusia}(\text{M}), \text{Hear\_alarm}(\text{M})\}$ . Here are some values of the corresponding coalgebra  $p_{al}: \text{At}_{al} \rightarrow \mathcal{M}_{pr}\mathcal{P}_f\text{At}_{al}$ :

$$p_{al}(\text{Hear\_alarm}(\text{M})) = 0.8\{\text{Alarm}, \text{Wake}(\text{M})\} + 0.3\{\text{Paracusia}(\text{M})\}$$

$$p_{al}(\text{Earthquake}) = 0.01\{\}$$

**Remark 3.2.** One might wonder why not simply adopt  $\mathcal{P}_f(\mathcal{P}_f(-) \times [0, 1])$  as the coalgebra type for PLP. The reason is that, although every ground PLP program generates a  $\mathcal{P}_f(\mathcal{P}_f(-) \times [0, 1])$ -coalgebra, such encoding fails to be a 1-1 correspondence: a clause  $\mathcal{C} \in \mathcal{P}_f(\text{At})$  may be associated with different values in  $[0, 1]$ , which violates the standard definition of PLP programs.

**3.2. Derivation Semantics.** In this section we are going to construct the final  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(-)$ -coalgebra, thus providing a semantic interpretation for probabilistic logic programs based on  $\text{At}$ .


Before the technical developments, we give an intuitive view on the semantics that the final coalgebra is going to provide. We shall represent each goal as a *stochastic derivation tree* in the final  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(-)$ -coalgebra. These trees are the probabilistic version of and-or derivation trees, which represent parallel SLD-resolutions for pure logic programming [GC94]. One can view a stochastic derivation trees as the unfolding of a goal under a PLP program.

**Definition 3.3** (Stochastic derivation trees). Given a ground PLP program  $\mathbb{P}$  based on  $\text{At}$ , and an atom  $A \in \text{At}$ , the *stochastic derivation tree* for  $A$  in  $\mathbb{P}$  is the possibly infinite tree  $\mathcal{T}$  such that:

- (1) Every node is either an atom-node (labelled with an atom  $A' \in \text{At}$ ) or a clause-node (labelled with  $\bullet$ ). These two types of nodes appear alternatingly in depth, in this order. In particular, the root is an atom-node labelled with  $A$ .
- (2) Each edge from an atom-node to its (clause-)children is labelled with a probability value.
- (3) Suppose  $s$  is an atom-node with label  $A'$ . Then for every clause  $r :: A' \leftarrow B_1, \dots, B_k$  in  $\mathbb{P}$ ,  $s$  has exactly one child  $t$  such that the edge  $s \rightarrow t$  is labelled with  $r$ , and  $t$  has exactly  $k$  children labelled with  $B_1, \dots, B_k$ , respectively.

The final coalgebra semantics  $\llbracket - \rrbracket_p$  for a program  $\mathbb{P}$  will map a goal  $A$  to the stochastic derivation tree representing all possible SLD-resolutions of  $A$  in  $\mathbb{P}$ .

**Example 3.4.** Continuing Example 2.1,  $\llbracket \text{Hear\_alarm}(\text{M}) \rrbracket_{p_{al}}$  is the stochastic derivation tree below. The subtree highlighted in red represents one of the successful proofs of  $\text{Hear\_alarm}(\text{M})$  in  $p_{al}$ : indeed, note that a single child is selected for each atom-node  $A$  (corresponding to a clause matching  $A$ ), all children of any clause-node are selected (corresponding to the atoms in the body of the clause), and the subtree has clause-nodes as leaves (all atoms are proven).



Any such subtree describes a proof, but does not yield a probability value to be associated to a goal — this is the remit of the distribution semantics, see Example 3.9 below.

In the remaining part of the section, we construct the cofree coalgebra for  $\mathcal{M}_{pr}\mathcal{P}_f$  via a so-called terminal sequence [Wor99], and obtain  $\llbracket - \rrbracket_p$  from the resulting universal property. We report the steps of the terminal sequence as they are instrumental in showing that the elements of the cofree coalgebra can be seen as stochastic derivation trees.

**Construction 3.5.** The terminal sequence for the functor  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(-) : \mathbf{Sets} \rightarrow \mathbf{Sets}$  consists of sequences of objects  $\{X_\alpha\}_{\alpha \in \mathbf{Ord}}$  and arrows  $\{\delta_\beta^\alpha : X_\alpha \rightarrow X_\beta\}_{\beta < \alpha \in \mathbf{Ord}}$  constructed by the following transfinite induction:

$$X_\alpha := \begin{cases} \text{At} & \alpha = 0 \\ \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\xi) & \alpha = \xi + 1 \\ \lim\{\delta_\xi^\chi \mid \xi < \chi < \alpha\} & \alpha \text{ is limit} \end{cases}$$

$$\delta_\beta^\alpha := \begin{cases} \pi_1 & \alpha = 1, \beta = 0 \\ id_{\text{At}} \times \mathcal{M}_{pr}\mathcal{P}_f(\delta_\xi^{\xi+1}) & \alpha = \beta + 1 = \xi + 2 \\ \text{the limit projections} & \alpha \text{ is limit, } \beta < \alpha \\ \text{universal map to } X_\beta & \beta \text{ is limit, } \alpha = \beta + 1 \end{cases}$$

**Proposition 3.6.** *The terminal sequence for the functor  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(-)$  converges to a limit  $X_\gamma$  such that  $X_\gamma \cong X_{\gamma+1}$ .*

*Proof.* We apply the following result from [Wor99]:

**Proposition 3.7** [Wor99, Corollary 3.3]. *If  $T$  is an accessible endofunctor on a locally presentable category, and if  $T$  preserves monics, then the terminal  $T$ -sequence  $\{A_\alpha, f_\beta^\alpha\}$  converges, necessarily to a terminal  $T$ -coalgebra.*

Since  $\mathbf{Sets}$  is a locally presentable category, it remains to verify the two conditions for  $\mathcal{M}_{pr}\mathcal{P}_f$ . It is well-known that  $\mathcal{P}_f$  is  $\omega$ -accessible, and  $\mathcal{M}_{pr}$  has the same property, see e.g. [Sil10, Prop. 6.1.2]. Because accessibility is defined in terms of colimit preservation, it is clearly preserved by composition, and thus  $\mathcal{M}_{pr}\mathcal{P}_f$  is also accessible. It remains to check that it preserves monics. For  $\mathcal{M}_{pr}$ , given any monomorphism  $i : C \rightarrow D$  in  $\mathbf{Sets}$ , suppose  $\mathcal{M}_{pr}(i)(\varphi) = \mathcal{M}_{pr}(i)(\varphi')$  for some  $\varphi, \varphi' \in \mathcal{M}_{pr}(C)$ . Then for any  $d \in D$ ,

$\mathcal{M}_{pr}(i)(\varphi)(d) = \mathcal{M}_{pr}(i)(\varphi')(d)$ . If we focus on the image  $i[C]$ , then there is an inverse function  $i^{-1} : i[C] \rightarrow C$ , and  $\mathcal{M}_{pr}(i)(\varphi) = \mathcal{M}_{pr}(i)(\varphi')$  implies that  $\varphi(i^{-1}(d)) = \varphi'(i^{-1}(d))$  for any  $d \in i[C]$ . But this simply means that  $\varphi = \varphi'$ . As the same is true for  $\mathcal{P}_f$  and the property is preserved by composition, we have that  $\mathcal{M}_{pr}\mathcal{P}_f$  preserves monics. Therefore we can conclude that the terminal sequence for  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f$  converges to the cofree  $\mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra on  $\text{At}$ .  $\square$

Note that  $X_{\gamma+1}$  is defined as  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\gamma)$ , and the above isomorphism makes  $X_\gamma \rightarrow \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\gamma)$  the final  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra — or, in other words, *cofree  $\mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra* on  $\text{At}$ . As for the tree representation of the elements of  $X_\gamma$ , recall that elements of the cofree  $\mathcal{P}_f\mathcal{P}_f$ -coalgebra on  $\text{At}$  can be seen as and-or trees [KMP10]. By replacing the first  $\mathcal{P}_f$  with  $\mathcal{M}_{pr}$ , effectively one adds probability values to the edges from and-nodes to or-nodes (which are edges from atom-nodes to clause-nodes in our stochastic derivation trees), as in (3.1). Thus stochastic derivation trees as in Definition 3.3 are elements of  $X_\gamma$ . The action of the coalgebra map  $\cong : X_\gamma \rightarrow \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\gamma)$  is best seen with an example: the tree  $\mathcal{T}$  in (3.1) (as an element of  $X_\gamma$ ) is mapped to the pair  $\langle \text{Hear\_alarm}(\text{M}), \varphi \rangle$ , where  $\varphi$  is the function  $\mathcal{P}_f(X_\gamma) \rightarrow [0, 1]$  assigning 0.8 to the set consisting of the subtrees of  $\mathcal{T}$  with root  $\text{Alarm}$  and with root  $\text{Wake}(\text{M})$ , 0.3 to the singleton consisting to the subtree of  $\mathcal{T}$  with root  $\text{Paracusia}(\text{M})$ , and 0 to any other finite set of trees.

With all the definitions at hand, it is straightforward to check that  $\llbracket - \rrbracket_p$  mapping  $A \in \text{At}$  to its stochastic derivation tree in  $\mathbb{P}$  makes the following diagram commute

$$\begin{array}{ccc}
 \text{At} & \xrightarrow{\llbracket - \rrbracket_p} & X_\gamma \\
 \downarrow \langle id, p \rangle & & \downarrow \cong \\
 \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(\text{At}) & \xrightarrow{id \times \mathcal{M}_{pr}\mathcal{P}_f(\llbracket - \rrbracket_p)} & \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\gamma)
 \end{array}$$

and thus by uniqueness it coincides with the  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra map provided by the universal property of the final  $\text{At} \times \mathcal{M}_{pr}\mathcal{P}_f$ -coalgebra  $X_\gamma \rightarrow \text{At} \times \mathcal{M}_{pr}\mathcal{P}_f(X_\gamma)$ .

**3.3. Distribution Semantics.** This section gives a coalgebraic representation of the usual *distribution semantics* of probabilistic logic programming. As in the previous section, before the technical developments we gather some preliminary intuition. Recall from Section 2 that the core of the distribution semantics is the probability distribution over the sub-programs (subsets of clauses) of a given program  $\mathbb{P}$ . These sub-programs are also called (possible) worlds, and the distribution semantics of a goal is the sum of the probabilities of all the worlds in which it is provable.

In order to encode this information as elements of a final coalgebra, we need to present it in tree-shape. Roughly speaking, we form a distribution over the sub-programs along the execution tree. This justifies the following notion of *distribution trees*.

**Definition 3.8** (Distribution trees). Given a PLP program  $\mathbb{P}$  and an atom  $A$ , the *distribution tree* for  $A$  in  $\mathbb{P}$  is the possibly infinite tree  $\mathcal{T}$  satisfying the following properties:


- (1) Every node is exactly one of the three kinds: atom-node (labelled with an atom  $A' \in \text{At}$ ), world-node (labelled with  $\circ$ ), clause-node (labelled with  $\bullet$ ). They appear alternately in this order in depth. In particular the root is an atom-node labelled with  $A$ .
- (2) Every edge from an atom-node to its (world-)children is labelled with a probability value, and the labels on all the edges (if exist) from the same atom-node sum up to 1.



- (3) Suppose  $s$  is an atom-node labelled with  $A'$ , and  $C = \{C_1, \dots, C_m\}$  is the set of all the clauses in  $\mathbb{P}$  whose head is  $A'$ . Then  $s$  has  $2^m$  (world-)children, each standing for a subset  $X$  of  $C$ . If a child  $t$  stands for  $X$ , then the edge  $s \rightarrow t$  is labelled with probability  $\prod_{C \in X} \text{Label}(C) \cdot \prod_{C' \in C \setminus X} (1 - \text{Label}(C'))$  — recall that  $\text{Label}(C)$  is the probability labelling  $C$ . Also,  $t$  has exactly  $|X|$ -many (clause-)children, each standing for a clause  $C \in X$ . If a child  $u$  stands for  $C = r :: A' \leftarrow B_1, \dots, B_k$ , then  $u$  has  $k$  (atom-)children, labelled with  $B_1, \dots, B_k$  respectively.

Comparing distribution trees with stochastic derivation trees (Definition 3.3), one can observe the addition of another class of nodes, representing possible worlds. Moreover, the possible worlds associated with an atom-node must form a probability distribution — as opposed to stochastic derivation trees, in which probabilities labelling parallel edges do not need to share any relationship. An example of the distribution tree associated with a goal is provided in the continuation of our leading example (Examples 2.1 and 3.4).

**Example 3.9.** In the context of Example 2.1, the distribution tree of  $\text{Hear\_alarm}(M)$  is depicted below, where we use grey shades to emphasise sets of edges expressing a probability distribution. Also, note the  $\circ$ s with no children, standing for empty worlds (namely worlds containing no clause).



In the literature, the distribution semantics usually associates with a goal a single probability value (2.1), rather than a whole tree. However, given the distribution tree it is straightforward to compute such probability. The subtree highlighted in red above describes a refutation of  $\text{Hear\_alarm}(M)$  with probability 0.000001296 (= the product of all the probabilities in the subtree). The sum of all the probabilities associated to such ‘proof’ subtrees yields the usual distribution semantics (2.1) — the computation is shown in detail in Appendix A.

In the remainder of this section, we focus on the coalgebraic characterisation of distribution trees and the associated semantics map. Our strategy will be to introduce a novel coalgebra type  $\mathcal{D}_{\leq 1} \mathcal{P}_f \mathcal{P}_f$ , such that distribution trees can be seen as elements of the cofree coalgebra. Then, we will provide a natural transformation  $\mathcal{M}_{pr} \Rightarrow \mathcal{D}_{\leq 1} \mathcal{P}_f$ , which can be used to transform stochastic derivation trees into distribution trees. Finally, composing the universal properties of these cofree coalgebras will yield the desired distribution semantics.

We begin with the definition of  $\mathcal{D}_{\leq 1} \mathcal{P}_f$ . This is simply the composite  $\mathcal{D}_{\leq 1} \mathcal{P}_f : \mathbf{Sets} \rightarrow \mathbf{Sets}$ , where  $\mathcal{D}_{\leq 1}$  is the *sub-probability distribution* functor. Recall that  $\mathcal{D}_{\leq 1}$  maps  $X$  to the

set of sub-probability distributions with finite supports on  $X$  (i.e., convex combinations of elements of  $X$  whose sum is less or equal to 1), and acts component-wise on functions.

**Remark 3.10.** Note that we cannot work with full probabilities  $\mathcal{D}$  here, since a goal may not match any clause. In such a case there is no world in which the goal is provable and its probability in the program is 0.

The next step is to recover distribution trees as the elements of the  $\mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$ -cofree coalgebra on  $\mathbf{At}$ . This goes via a terminal sequence, similarly to the case of  $\mathcal{M}_{pr}\mathcal{P}_f$  in the previous section. We will not go into full detail, but only mention that the terminal sequence for  $\mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f(-) : \mathbf{Sets} \rightarrow \mathbf{Sets}$  is constructed as the one for  $\mathbf{At} \times \mathcal{M}_{pr}\mathcal{P}_f(-) : \mathbf{Sets} \rightarrow \mathbf{Sets}$  (Construction 3.5), with  $\mathcal{D}_{\leq 1}\mathcal{P}_f$  replacing  $\mathcal{M}_{pr}$ .

**Proposition 3.11.** *The terminal sequence of  $\mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f(-)$  converges at some limit ordinal  $\chi$ , and  $(\lambda_\chi^{X+1})^{-1} : Y_\chi \rightarrow \mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f Y_\chi$  is the final  $\mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$  coalgebra.*

*Proof.* As for Proposition 3.6, by [Wor99, Cor. 3.3] it suffices to show that  $\mathcal{D}_{\leq 1}\mathcal{P}_f$  is accessible and preserves monos. Both are simple exercises; in particular, see [BSdV04] for accessibility of  $\mathcal{D}_{\leq 1}$ .  $\square$

The association of distribution trees with elements of  $Y_\chi$  is suggested by the type  $\mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f$ . Indeed,  $\mathbf{At} \times \mathcal{D}_{\leq 1}$  is the layer of atom-nodes, labelled with elements of  $\mathbf{At}$  and with outgoing edges forming a sub-probability distribution; the first  $\mathcal{P}_f$  is the layer of world-nodes; the second  $\mathcal{P}_f$  is the layer of clause-nodes. The coalgebra map  $Y_\chi \rightarrow \mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f Y_\chi$  associates a goal to subtrees of its distribution trees, analogously to the coalgebra structure on stochastic derivation trees in the previous section.

The last ingredient we need is a translation of stochastic derivation trees into distribution trees. We formalise this as a natural transformation  $\mathbf{pw} : \mathcal{M}_{pr} \Rightarrow \mathcal{D}_{\leq 1}\mathcal{P}_f$ . The naturality of  $\mathbf{pw}$  can be checked with a simple calculation.

**Definition 3.12.** The ‘possible worlds’ natural transformation  $\mathbf{pw} : \mathcal{M}_{pr} \Rightarrow \mathcal{D}_{\leq 1}\mathcal{P}_f$  is defined by  $\mathbf{pw}_X : \varphi \mapsto \sum_{Y \subseteq \text{supp}(\varphi)} r_Y Y$ , where each  $r_Y = \prod_{y \in Y} \varphi(y) \cdot \prod_{y' \in \text{supp}(\varphi) \setminus Y} (1 - \varphi(y'))$ . In particular, when  $\text{supp}(\varphi)$  is empty,  $\mathbf{pw}_X(\varphi)$  is the empty sub-distribution  $\emptyset$ .

Now we have all the ingredients to characterise the distribution semantics coalgebraically, as the morphism  $\llbracket - \rrbracket_p : \mathbf{At} \rightarrow Y_\chi$  defined by the following diagram, which maps  $A \in \mathbf{At}$  to its distribution tree in  $p$ .

$$\begin{array}{ccccc}
 \mathbf{At} & \xrightarrow{\llbracket - \rrbracket_p} & X_\gamma & \xrightarrow{!} & Y_\chi \\
 \downarrow \langle id_{\mathbf{At}}, p \rangle & & \downarrow \cong & & \downarrow \cong \\
 \mathbf{At} \times \mathcal{M}_{pr}\mathcal{P}_f\mathbf{At} & \xrightarrow{id_{\mathbf{At}} \times \mathcal{M}_{pr}\mathcal{P}_f(\llbracket - \rrbracket_p)} & \mathbf{At} \times \mathcal{M}_{pr}\mathcal{P}_f X_\gamma & & \\
 \downarrow id_{\mathbf{At}} \times \mathbf{pw}_{\mathcal{P}_f(\mathbf{At})} & & \downarrow id_{\mathbf{At}} \times \mathbf{pw}_{\mathcal{P}_f(X_\gamma)} & & \\
 \mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f\mathbf{At} & \xrightarrow{id_{\mathbf{At}} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f(\llbracket - \rrbracket_p)} & \mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f X_\gamma & \xrightarrow{id_{\mathbf{At}} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f(!)} & \mathbf{At} \times \mathcal{D}_{\leq 1}\mathcal{P}_f\mathcal{P}_f Y_\chi \\
 & & & & (3.2)
 \end{array}$$

Note the use of `pw` to extend probabilistic logic programs and stochastic derivation trees to the same coalgebra type as distribution trees. Then the distribution semantics  $\langle\langle - \rangle\rangle_p$  is uniquely defined by the universal property of the final  $\text{At} \times \mathcal{D}_{\leq 1} \mathcal{P}_f \mathcal{P}_f$ -coalgebra. By uniqueness, it can also be computed as the composite  $! \circ \llbracket - \rrbracket_p$ , that is, first one derives the semantics  $\llbracket - \rrbracket_p$ , then applies the translation `pw` to each level of the resulting stochastic derivation tree, in order to turn it into a distribution tree.

#### 4. GENERAL CASE

We now generalise our coalgebraic treatment to arbitrary probabilistic logic programs and goals, possibly including variables. The section has a similar structure as the one devoted to the ground case. First, in Subsection 4.1, we give a coalgebraic representation for general PLP, and equip it with a final coalgebra semantics in terms of stochastic derivation trees (Subsection 4.2). Next, in Subsection 4.3, we study the coalgebraic representation of the distribution semantics. We begin by introducing our leading example — an extension of Example 2.1.


**Example 4.1.** We tweak the ground program of Example 2.1. Now it is not just Mary that may hear the alarm, but also her neighbours. There is a small probability that the alarm rings because someone passes too close to Mary’s house. However, we can only estimate the possibility of paracusia and being awake for Mary, not the neighbours. The revised program, which by abuse of notation we also call  $\mathbb{P}^{al}$ , is based on an extension of the language in Example 2.1: we add a new 1-ary function symbol  $\text{Neigh}^1$  to the signature  $\Sigma_{al}$ , and a new 1-ary predicate  $\text{PassBy}(-)$  to the alphabet. Note the appearance of a variable  $x$ .

0.01 :: Earthquake	←	0.5 :: Alarm	← Earthquake
0.2 :: Burglary	←	0.9 :: Alarm	← Burglary
0.6 :: Wake(Mary)	←	0.1 :: Alarm	← PassBy(x)
0.01 :: Paracusia(Mary)	←	0.3 :: Hear_alarm(x)	← Paracusia(x)
0.8 :: Wake(Neigh(x))	← Wake(x)	0.8 :: Hear_alarm(x)	← Alarm, Wake(x)

As we want to maintain our approach a direct generalisation of the coalgebraic semantics [BZ15] for pure logic programs, the derivation semantics  $\llbracket - \rrbracket$  for general PLP will represent resolution by *unification*. This means that, at each step of the computation, given a goal  $A$ , one seeks substitutions  $\theta, \tau$  such that  $A\theta = H\tau$  for the head  $H$  of some clause in the program. As a roadmap, we anticipate the way this computation is represented in terms of stochastic derivation trees (Definition 4.7 below), with a continuation of our leading example.

**Example 4.2.** In the context of Example 4.1, the tree for  $\llbracket \text{Hear\_alarm}(x) \rrbracket_{\mathbb{P}^{al}}$  is (partially) depicted below. Compared to the ground case (Example 3.4), now substitutions applied on the goal side appear explicitly as labels. Moreover, note that for each atom  $A\theta$  and clause  $\mathcal{C}$ , there might be multiple substitutions  $\tau$  such that  $A\theta = \text{Head}(\mathcal{C})\tau$ . This is presented in the following tree (4.1) by the  $\blacklozenge$ -labelled nodes, which are children of the clause-nodes (labelled

with  $\bullet$ ). We abbreviate Neigh as N and Mary as M.



Resolution by unification as above will be implemented in two stages. The first step is devising a map for term-matching. Assuming that the substitution instance  $A\theta$  of a goal  $A$  is already given, we define  $p$  performing term-matching of  $A\theta$  in a given program  $\mathbb{P}$ :

$$p(A\theta): \{B_1\tau_i, \dots, B_k\tau_i\}_{i \in I \subseteq \mathbb{N}} \mapsto \begin{cases} r & (r :: H \leftarrow B_1, \dots, B_k) \in \mathbb{P} \text{ and} \\ & I \text{ contains all } i \text{ s.t. } A\theta = H\tau_i \\ 0 & \text{otherwise.} \end{cases} \quad (4.2)$$

Intuitively, one application of such map is represented in a tree structure as Example 4.2 by the first two layers of the subtree rooted at  $\theta$ . The reason why the domain of  $p(A)$  is a countable set  $\{B_1\tau_i, \dots, B_k\tau_i\}_{i \in I \subseteq \mathbb{N}}$  of instances of the same body  $B_1, \dots, B_k$  is that the same clause may match a goal with countably many different substitutions  $\tau_i$ . For example in the bottom part of (4.1) there are countably infinite substitutions  $\tau_i$  matching the head of  $\text{Alarm} \leftarrow \text{PassBy}(x)$  to the goal  $\text{Alarm}$ , substituting  $x$  with  $\text{Mary}$ ,  $\text{Neigh}(\text{Mary})$ ,  $\text{Neigh}(x), \dots$ . This will be reflected in the coalgebraic representation of PLP (see (4.4) below) by the use of the countable powerset functor  $\mathcal{P}_c$ .

In order to model arbitrary unification, the second step is considering all substitutions  $\theta$  on the goal  $A$  such that a term-matcher for  $A\theta$  exists. There is an elegant categorical construction [BZ15] packing together these two steps into a single coalgebra map. We will start with this in the next session.

**4.1. Coalgebraic Representation of PLP.** Towards a categorification of general PLP, the first concern is to account for the presence of variables in atoms. This is standardly done by letting the space of atoms on an alphabet  $\mathcal{A}$  be a presheaf  $\text{At}: \mathbf{L}_\Sigma^{\text{op}} \rightarrow \mathbf{Sets}$  rather than a set. Here the index category  $\mathbf{L}_\Sigma^{\text{op}}$  is the opposite *Lawvere Theory* of  $\Sigma$  (see Section 2). For each  $n \in \text{Ob}(\mathbf{L}_\Sigma^{\text{op}})$ ,  $\text{At}(n)$  is defined as the set of  $\mathcal{A}$ -atoms in context  $n$ . Given an  $n$ -tuple  $\theta = \langle t_1, \dots, t_n \rangle \in \mathbf{L}_\Sigma^{\text{op}}[n, m]$  of  $\Sigma$ -terms in context  $m$ ,  $\text{At}(\theta): \text{At}(n) \rightarrow \text{At}(m)$  is defined by substitution, namely  $\text{At}(\theta)(A) = A\theta$ , for any  $A \in \text{At}(n)$ .

As observed in [KP11] for pure logic programs, if we naively try to model our specification (4.2) for  $p$  as a coalgebra on  $\text{At}$ , we run into problems: indeed  $p$  is not a natural transformation, thus not a morphism between presheaves. Intuitively, this is because the existence of a term-matching for a substitution instance  $A\theta$  of  $A$  does not necessarily imply the existence

of a term-matching for  $A$  itself. For pure logic programs, this problem can be solved in at least two ways. First, [KP11] relaxes naturality by changing the base category of presheaves from **Sets** to **Poset**. We take here the second route, namely give a ‘saturated’ coalgebraic treatment of PLP, generalising the modelling of pure logic programs proposed in [BZ15]. This approach has the advantage of letting us work with **Sets**-based presheaves, and be still able to recover term-matching via a ‘desaturation’ operation — see [BZ15] and Appendix B.

**The Saturation Adjunction.** To this aim, we briefly recall the saturated approach from [BZ15]. The central piece is the adjunction  $\mathcal{U} \dashv \mathcal{K}$  on presheaf categories, as on the left below.

$$\begin{array}{ccc}
 \mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}} & \begin{array}{c} \xrightarrow{\mathcal{U}} \\ \perp \\ \xleftarrow{\mathcal{K}} \end{array} & \mathbf{Sets}^{|\mathbf{L}_\Sigma^{\text{op}}|} \\
 & & \\
 & & \begin{array}{ccc} |\mathbf{L}_\Sigma^{\text{op}}| & \xrightarrow{\iota} & \mathbf{L}_\Sigma^{\text{op}} \\ \mathbf{F} \downarrow & \swarrow \mathcal{K}(\mathbf{F}) & \\ \mathbf{Sets} & & \end{array}
 \end{array} \quad (4.3)$$

Here  $|\mathbf{L}_\Sigma^{\text{op}}|$  is the discretisation of  $\mathbf{L}_\Sigma^{\text{op}}$ , i.e. all the arrows but the identities are dropped. The left adjoint  $\mathcal{U}$  is the forgetful functor, given by precomposition with the obvious inclusion  $\iota: |\mathbf{L}_\Sigma^{\text{op}}| \rightarrow \mathbf{L}_\Sigma^{\text{op}}$ .  $\mathcal{U}$  has a right adjoint  $\mathcal{K} = \text{Ran}\iota: \mathbf{Sets}^{|\mathbf{L}_\Sigma^{\text{op}}|} \rightarrow \mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}}$ , which sends every presheaf  $\mathbf{F}: |\mathbf{L}_\Sigma^{\text{op}}| \rightarrow \mathbf{Sets}$  to its *right Kan extension* along  $\iota$ , as in the rightmost diagram in (4.3). The definition of  $\mathcal{K}$  can be computed [Mac71] as follows:

- On objects  $\mathbf{F} \in \text{Ob}(\mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}})$ , the presheaf  $\mathcal{K}(\mathbf{F}): \mathbf{L}_\Sigma^{\text{op}} \rightarrow \mathbf{Sets}$  is defined by letting  $\mathcal{K}(\mathbf{F})(n)$  be the product  $\prod_{\theta \in \mathbf{L}_\Sigma^{\text{op}}[n,m]} \mathbf{F}(m)$ , where  $m$  ranges over  $\text{Ob}(\mathbf{L}_\Sigma^{\text{op}})$ . Intuitively, every element in  $\mathcal{K}(\mathbf{F})(n)$  is a tuple with index set  $\bigcup_{m \in \text{Ob}(\mathbf{L}_\Sigma^{\text{op}})} \mathbf{L}_\Sigma^{\text{op}}[n,m]$ , and its component at index  $\theta: n \rightarrow m$  must be an element in  $\mathbf{F}(m)$ . We follow the convention of [BZ15] and write  $\dot{x}, \dot{y}, \dots$  for such tuples, and  $\dot{x}(\theta)$  for the component of  $\dot{x}$  at index  $\theta$ .

With this convention, given an arrow  $\sigma \in \mathbf{L}_\Sigma^{\text{op}}[n, n']$ ,  $\mathcal{K}(\mathbf{F})(\sigma)$  is defined by pointwise substitution as the mapping of the tuple  $\dot{x}$  to the tuple  $\langle \dot{x}(\theta \circ \sigma) \rangle_{\theta: n' \rightarrow m}$ .

- On arrows, given a morphism  $\alpha: \mathbf{F} \rightarrow \mathbf{G}$  in  $\mathbf{Sets}^{|\mathbf{L}_\Sigma^{\text{op}}|}$ ,  $\mathcal{K}(\alpha)$  is a natural transformation  $\mathcal{K}(\mathbf{F}) \rightarrow \mathcal{K}(\mathbf{G})$  defined pointwisely as  $\mathcal{K}(\alpha)(n): \dot{x} \mapsto \langle \alpha_m(\dot{x}(\theta)) \rangle_{\theta: n \rightarrow m}$ .

It is also useful to record the unit  $\eta: 1 \rightarrow \mathcal{K}\mathcal{U}$  of the adjunction  $\mathcal{U} \dashv \mathcal{K}$ . Given a presheaf  $\mathbf{F}: \mathbf{L}_\Sigma^{\text{op}} \rightarrow \mathbf{Sets}$ ,  $\eta_{\mathbf{F}}: \mathbf{F} \rightarrow \mathcal{K}\mathcal{U}\mathbf{F}$  is a natural transformation defined by  $\eta_{\mathbf{F}}(n): x \mapsto \langle \mathbf{F}(\theta)(x) \rangle_{\theta: n \rightarrow m}$ .

**Saturation in PLP.** We now construct the coalgebra structure on the presheaf **At** modelling PLP. First, we are now able to represent  $p$  in (4.2) as a coalgebra map. The aforementioned naturality issue is solved by defining it as a morphism in  $\mathbf{Sets}^{|\mathbf{L}_\Sigma^{\text{op}}|}$  rather than in  $\mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}}$ , thus making naturality trivial. The coalgebra  $p$  will have the following type

$$p: \mathcal{U}\text{At} \rightarrow \widehat{\mathcal{M}}_{pr} \widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}\text{At} \quad (4.4)$$

where  $\widehat{(\cdot)}$  is the obvious extension of **Sets**-endofunctors to  $\mathbf{Sets}^{|\mathbf{L}_\Sigma^{\text{op}}|}$ -endofunctors, defined by functor precomposition. With respect to the ground case, note the insertion of  $\widehat{\mathcal{P}}_c$ , the lifting of the *countable* powerset functor, in order to account for the countably many instances of a clause that may match the given goal (*cf.* the discussion below (4.2)).

**Example 4.3.** Our program  $\mathbb{P}^{al}$  (Example 4.1) is based on  $\text{At}_{al}: \mathbf{L}_{\Sigma_{al}}^{\text{op}} \rightarrow \mathbf{Sets}$ . Some of its values are  $\text{At}_{al}(0) = \{\text{Mary}, \text{Neigh}(\text{Mary}), \text{Neigh}(\text{Neigh}(\text{Mary})), \dots\}$  and  $\text{At}_{al}(1) = \{x, \text{Mary}, \text{Neigh}(x), \text{Neigh}(\text{Mary}), \dots\}$ . Part of the coalgebra  $p_{al}$  modelling the program  $\mathbb{P}^{al}$  is as follows (*cf.* the tree (4.1)). For space reason, we abbreviate **Mary** as **M** in the second

equation.

$$\begin{aligned} (p_{al})_0(\text{Hear\_alarm}(\text{Mary})) &= 0.8\{\{\text{Alarm}, \text{Wake}(\text{Mary})\}\} + 0.3\{\{\text{Parasusia}(\text{Mary})\}\} \\ (p_{al})_1(\text{Alarm}) &= 0.5\{\{\text{Earthquake}\}\} + 0.9\{\{\text{Burglary}\}\} \\ &\quad + 0.1\{\{\text{PassBy}(\text{M})\}, \{\text{PassBy}(\text{Neigh}(\text{M}))\}, \{\text{PassBy}(\text{Neigh}(\text{x}))\}, \dots\} \end{aligned}$$

The universal property of the adjunction (4.3) gives a canonical ‘lifting’ of  $p$  to a  $\mathcal{K}\widehat{\mathcal{M}}_{pr}\widehat{\mathcal{P}}_c\widehat{\mathcal{P}}_f\mathcal{U}$ -coalgebra  $p^\sharp$  on  $\text{At}$ , performing unification rather than just term-matching:

$$p^\sharp := \text{At} \xrightarrow{\eta_{\text{At}}} \mathcal{K}\mathcal{U}\text{At} \xrightarrow{\mathcal{K}p} \mathcal{K}\widehat{\mathcal{M}}_{pr}\widehat{\mathcal{P}}_c\widehat{\mathcal{P}}_f\mathcal{U}\text{At} \quad (4.5)$$

where  $\eta$  is the unit of the adjunction, as defined above. Spelling it out,  $p^\sharp$  is the mapping

$$p_n^\sharp : A \in \text{At}(n) \mapsto \langle p_m(A\theta) \rangle_{\theta: n \rightarrow m}$$

Intuitively,  $p_n^\sharp$  retrieves all the unifiers  $\langle \theta, \tau \rangle$  of  $A$  and head  $H$  in  $\mathbb{P}$ : first, we have  $A\theta \in \text{At}(m)$  as a component of the saturation of  $A$  by  $\eta_{\text{At}}$ ; then we term-match  $H$  with  $A\theta$  by  $\mathcal{K}p_m$ .

**Remark 4.4.** Note that the parameter  $n \in \text{Ob}(\mathbf{L}_\Sigma^{\text{op}})$  in the natural transformation  $p^\sharp$  fixes the pool  $\{x_1, \dots, x_n\}$  of variables appearing in the atoms (and relative substitutions) that are considered in the computation. Analogously to the case of pure logic programs [KP11, BZ15], it is intended that such  $n$  can always be chosen ‘big enough’ so that all the relevant substitution instances of the current goal and clauses in the program are covered — note the variables occurring therein always form a *finite* set, included in  $\{x_1, \dots, x_m\}$  for some  $m \in \mathbb{N}$ .

**4.2. Derivation Semantics.** Once we have identified our coalgebra type, the construction leading to the derivation semantics  $\llbracket - \rrbracket_{p^\sharp}$  for general PLP is completely analogous to the ground case. One can define the cofree coalgebra for  $\mathcal{K}\widehat{\mathcal{M}}_{pr}\widehat{\mathcal{P}}_c\widehat{\mathcal{P}}_f\mathcal{U}(-)$  by terminal sequence, similarly to Construction 3.5. For simplicity, henceforth we denote the functor  $\mathcal{K}\widehat{\mathcal{M}}_{pr}\widehat{\mathcal{P}}_c\widehat{\mathcal{P}}_f\mathcal{U}(-)$  by  $\mathcal{S}$ .

**Construction 4.5.** The terminal sequence for  $\text{At} \times \mathcal{S}(-) : \mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}} \rightarrow \mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}}$  consists of a sequence of objects  $X_\alpha$  and morphisms  $\delta_\alpha^\beta : X_\beta \rightarrow X_\alpha$ , for  $\alpha < \beta \in \mathbf{Ord}$ , defined analogously to Construction 3.5, with  $p^\sharp$  and  $\mathcal{S}$  replacing  $p$  and  $\mathcal{M}_{pr}\mathcal{P}_f$ .

This terminal sequence converges by the following proposition.

**Proposition 4.6.**  $\mathcal{S}$  is accessible, and preserves monomorphisms.

*Proof.* Since both properties are preserved by composition, it suffices to show that they hold for all the component functors. For  $\widehat{\mathcal{M}}_{pr}$ ,  $\widehat{\mathcal{P}}_c$  and  $\widehat{\mathcal{P}}_f$ , they follow from accessibility and mono-preservation of  $\mathcal{M}_{pr}$ ,  $\mathcal{P}_c$  and  $\mathcal{P}_f$  (see Proposition 3.6), as (co)limits in presheaf categories are computed pointwise. For  $\mathcal{K}$  and  $\mathcal{U}$ , these properties are proven in [BZ15].  $\square$

Therefore the terminal sequence for  $\text{At} \times \mathcal{S}(-)$  converges at some limit ordinal, say  $\gamma$ , yielding the final  $\text{At} \times \mathcal{S}(-)$ -coalgebra  $X_\gamma \xrightarrow{\cong} \text{At} \times \mathcal{S}(X_\gamma)$ . The derivation semantics is then

defined  $\llbracket - \rrbracket_{\mathbb{P}^\#} : \text{At} \rightarrow X_\gamma$  by universal property as in the diagram below.

$$\begin{array}{ccc}
 \text{At} & \xrightarrow{\llbracket - \rrbracket_{\mathbb{P}^\#}} & X_\gamma \\
 \downarrow \langle id_{\text{At}}, p^\# \rangle & & \downarrow \cong \\
 \text{At} \times \mathcal{S}(\text{At}) & \xrightarrow{id_{\text{At}} \times \llbracket - \rrbracket_{\mathbb{P}^\#}} & \text{At} \times \mathcal{S}(X_\gamma)
 \end{array} \tag{4.6}$$

A careful inspection of the terminal sequence constructing  $X_\gamma$  allows to infer a representation of its elements as trees, among which we have those representing computations by unification of goals in a PLP program. We call these *stochastic saturated derivation trees*, as they extend the derivation trees of Definition 3.3 and are the probabilistic variant of saturated and-or trees in [BZ15]. Using (4.6) one can easily verify that  $\llbracket A \rrbracket_{\mathbb{P}^\#}$  is indeed the stochastic saturated derivation tree for a given goal  $A$ . Example 4.2 provides a pictorial representation of one such tree.

**Definition 4.7** (Stochastic saturated derivation trees). Given a probabilistic logic program  $\mathbb{P}$ , a natural number  $n$  and an atom  $A \in \text{At}(n)$ . The *stochastic saturated derivation tree* for  $A$  in  $\mathbb{P}$  is the possibly infinite tree  $\mathcal{T}$  satisfying the following properties:

- (1) There are four kinds of nodes: atom-node (labelled with an atom), substitution-node (labelled with a substitution), clause-node (labelled with  $\bullet$ ), instance-node (labelled with  $\blacklozenge$ ), appearing alternatively in depth in this order. The root is an atom-node with label  $A$ .
- (2) Each edge between a substitution node and its clause-node children is labelled with some  $r \in [0, 1]$ .
- (3) Suppose an atom-node  $s$  is labelled with  $A' \in \text{At}(n')$ . For every substitution  $\theta: n' \rightarrow m'$ ,  $s$  has exactly one (substitution-node) child  $t$  labelled with  $\theta$ . For every clause  $r :: H \leftarrow B_1, \dots, B_k$  in  $\mathbb{P}$  such that  $H$  matches  $A'\theta$  (via some substitution),  $t$  has exactly one (clause-)child  $u$ , and edge  $t \rightarrow u$  is labelled with  $r$ . Then for every substitution  $\tau$  such that  $A'\theta = H\tau$  and  $B_1\tau, \dots, B_k\tau \in \text{At}(m')$ ,  $u$  has exactly one (instance-)child  $v$ . Also  $v$  has exactly  $|\{B_1, \dots, B_k\}|_\tau$ -many (atom-)children, each labelled with one element in  $\{B_1, \dots, B_k\}_\tau$ .

**4.3. Distribution Semantics.** In this section we conclude by giving a coalgebraic perspective on the distribution semantics  $\langle\langle - \rangle\rangle$  for general PLP. Mimicking the ground case (Section 3.3), this will be presented as an extension of the derivation semantics, via a ‘possible worlds’ natural transformation. Also in the general case, we want to guarantee that a single probability value is computable for a given goal  $A$  from the corresponding tree  $\langle\langle A \rangle\rangle$  in the final coalgebra — whenever this probability is also computable in the ‘traditional’ way (see (2.1)) of giving distribution semantics to PLP. In this respect, the presence of variables and substitutions poses additional challenges, for which we refer to Appendices A and B. In a nutshell, the issue is that the distribution semantics counts the existence of a clause in the program rather than the number of times a clause is used in the computation. This means that every clause is counted at most once, independently from how many times that clause is used again in the computation. Note that our tree representation, as in the ground case, presents the resolution (thus computation aspect) of the program. So in our tree representation we need to give enough information to determine which clause is used at

each step of the computation, so that a second use can be easily detected. However, neither our saturated derivation trees, nor a ‘naive’ extension of them to distribution trees, carry such information: what appears in there is only the instantiated heads and bodies, but in general one cannot retrieve  $A$  from a substitution  $\theta$  and the instantiation  $A\theta$ . This is best illustrated via a simple example.

**Example 4.8.** Consider the following program, based on the signature  $\Sigma = \{a^0\}$  and two 1-ary predicates  $P, Q$ . It consists of two clauses:

$$0.5 :: P(x_1) \leftarrow Q(x_1) \quad | \quad 0.5 :: P(x_1) \leftarrow Q(x_2)$$

The goal  $P(a)$  matches the head of both clauses. However, given the sole information of the next goal being  $Q(a)$ , it is impossible to say whether the first clause has been used, instantiated with  $x_1 \mapsto a$ , or the second clause has been used, instantiated with  $x_1 \mapsto a, x_2 \mapsto a$ .

This observation motivates, as intermediate step towards the distribution semantics, the addition of labels to clause-nodes in derivation trees, in order to make explicit which clause is being applied. From the coalgebraic viewpoint, this just amounts to an extension of the type of the term-matching coalgebra:

$$\tilde{p}: \mathcal{UAt} \rightarrow \widehat{\mathcal{M}}_{pr}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{UAt} \times (\mathcal{UAt} \times \widehat{\mathcal{P}}_f \mathcal{UAt})) \quad (4.7)$$

Note the insertion of  $(-) \times (\mathcal{UAt} \times \widehat{\mathcal{P}}_f \mathcal{UAt})$ , which allows us to indicate at each step the head ( $\mathcal{UAt}$ ) and the body ( $\widehat{\mathcal{P}}_f \mathcal{UAt}$ ) of the clause being used, its probability label being already given by  $\widehat{\mathcal{M}}_{pr}$ . More formally, for any  $n$  and atom  $A \in \text{At}(n)$ , we define<sup>1</sup>

$$\tilde{p}_n(A): \langle \{B_1\tau_i, \dots, B_k\tau_i\}_{i \in \mathbb{I} \subseteq \mathbb{N}}, \langle H, \{B_1, \dots, B_k\} \rangle \rangle \mapsto \begin{cases} r & (r :: H \leftarrow B_1, \dots, B_k) \in \mathbb{P} \\ & \text{and } H\tau_i = A \\ 0 & \text{otherwise} \end{cases}$$

As in the case of  $p$  in (4.4), we can move from term-matching to unification by using the universal property of the adjunction  $\mathcal{U} \dashv \mathcal{K}$ , yielding  $\tilde{p}^\sharp: \text{At} \rightarrow \mathcal{K}\widehat{\mathcal{M}}_{pr}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{UAt} \times (\mathcal{UAt} \times \widehat{\mathcal{P}}_f \mathcal{UAt}))$ . For simplicity henceforth we denote the functor  $\mathcal{K}\widehat{\mathcal{M}}_{pr}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}(-) \times (\mathcal{UAt} \times \widehat{\mathcal{P}}_f \mathcal{UAt}))$  by  $\mathcal{R}$ .

We are now able to conclude our characterisation of the distribution semantics. The ‘possible worlds’ transformation  $\text{pw}: \mathcal{M}_{pr} \Rightarrow \mathcal{D}_{\leq 1} \mathcal{P}_f$  (Definition 3.12) yields a natural transformation  $\widehat{\text{pw}}: \widehat{\mathcal{M}}_{pr} \rightarrow \widehat{\mathcal{D}}_{\leq 1} \widehat{\mathcal{P}}_f$ , defined pointwise by  $\text{pw}$ . We can use  $\widehat{\text{pw}}$  to translate  $\mathcal{R}$  into the functor  $\mathcal{K}\widehat{\mathcal{D}}_{\leq 1} \widehat{\mathcal{P}}_f(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}(-) \times (\mathcal{UAt} \times \widehat{\mathcal{P}}_f \mathcal{UAt}))$ , abbreviated as  $\mathcal{O}$ , which is going to give the type of saturated distribution trees for general PLP programs.

As a simple extension of the developments in Section 4.2, we can construct the cofree  $\mathcal{R}$ -coalgebra  $\Phi \xrightarrow{\cong} \text{At} \times \mathcal{R}(\Phi)$  via a terminal sequence. Similarly, one can obtain the cofree  $\mathcal{O}$ -coalgebra  $\Psi \xrightarrow{\cong} \text{At} \times \mathcal{O}(\Psi)$ . By the universal property of  $\Psi$ , all these ingredients get

<sup>1</sup>As noted in Remark 4.4, instantiating  $\tilde{p}$  to some  $n \in \text{Ob}(\mathbf{L}_\Sigma^{\text{op}})$  fixes a variable context  $\{x_1, \dots, x_n\}$  both for the goal and the clause labels. In practice, because the set of clauses is always finite, it suffices to chose  $n$  ‘big enough’ so that the variables appearing in the clauses are included in  $\{x_1, \dots, x_n\}$ .



together in the definition of the distribution semantics  $\langle\langle - \rangle\rangle_{\tilde{p}^\#}$  for arbitrary PLP programs  $\tilde{p}^\#$

$$\begin{array}{ccccc}
 & & \langle\langle - \rangle\rangle_{\tilde{p}^\#} & & \\
 & \text{At} & \xrightarrow{\quad !_\Phi \quad} & \Phi & \xrightarrow{\quad !_\Psi \quad} & \Psi \\
 & \downarrow \langle id_{\text{At}}, \tilde{p}^\# \rangle & & \downarrow \cong & & \downarrow \cong \\
 \text{At} \times \mathcal{R}\text{At} & \xrightarrow{id_{\text{At}} \times \mathcal{R}(!_\Phi)} & & \text{At} \times \mathcal{R}\Phi & & \\
 \downarrow id_{\text{At}} \times \mathcal{K}\widehat{pw} & & & \downarrow id_{\text{At}} \times \mathcal{K}\widehat{pw} & & \\
 \text{At} \times \mathcal{O}\text{At} & \xrightarrow{id_{\text{At}} \times \mathcal{O}(!_\Phi)} & & \text{At} \times \mathcal{O}\Phi & \xrightarrow{id_{\text{At}} \times \mathcal{O}(!_\Psi)} & \text{At} \times \mathcal{O}\Psi
 \end{array}$$

where  $!_\Phi$  and  $!_\Psi$  are given by the evident universal properties, and show the role of the cofree  $\mathcal{R}$ -coalgebra  $\Phi$  as an intermediate step. The layered construction of final coalgebras  $\Psi$  and  $\Phi$ , together with the above characterisation of  $\langle\langle - \rangle\rangle_{\tilde{p}^\#}$ , allow to conclude that the distribution semantics for the program  $\tilde{p}^\#$  maps a goal  $A$  to its *saturated distribution tree*  $\langle\langle A \rangle\rangle_{\tilde{p}^\#}$ , as formally defined below.

**Definition 4.9** (Saturated distribution tree). The *saturated distribution tree* for  $A \in \text{At}(n)$  in  $\mathbb{P}$  is the possibly infinite  $\mathcal{T}$  satisfying the following properties based on Definition 4.7:

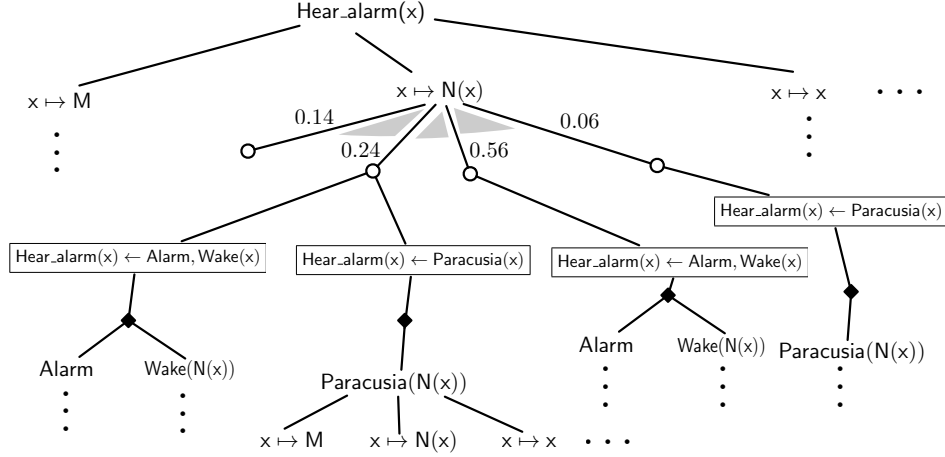
- (1) There are five kinds of nodes: in addition to the atom-, substitution-, clause- and instance-nodes, there are world-nodes. The world-nodes are children of the substitution-nodes, and parents of the clause nodes. The root and the order of the rest of the nodes are the same as in Definition 4.7, condition **1**. The clause-nodes are now labelled with clauses in  $\mathbb{P}$ .
- (2) Suppose  $s$  is an atom-node labelled with  $A' \in \text{At}(n')$ , and  $t$  is a substitution-child of  $s$  labelled with  $\theta: n' \rightarrow m$ . Let  $C$  be the set of all clauses  $\mathcal{C}$  such that  $\text{Head}(\mathcal{C})$  matches  $A'\theta$ . Then  $t$  has  $2^{|C|}$  world-children, each representing a subset  $X$  of  $C$ . If a world-child  $u$  of  $t$  represents subset  $X$ , then the edge  $t \rightarrow u$  has probability label  $\prod_{\mathcal{C} \in X} \text{Label}(\mathcal{C}) \cdot \prod_{\mathcal{C}' \in C \setminus X} (1 - \text{Label}(\mathcal{C}'))$ . Also  $u$  has  $|X|$  clause-children, one for each clause  $\mathcal{C} \in X$ , labelled with the corresponding clause. The rest for clause-nodes and instance-nodes are the same as in Definition 4.7, condition **3**.

**Remark 4.10.** Note that, in principle, saturated distribution trees could be defined coalgebraically without the intermediate step of adding clause labels. This is to be expected: coalgebra typically captures the one-step, ‘local’ behaviour of a system. On the other hand, as explained, the need for clause labels is dictated by a computational aspect involving the depth of distribution trees, that is, a ‘non-local’ dimension of the system.

We conclude with the pictorial representation of the saturated distribution tree of a goal in our leading example.

**Example 4.11.** In the context of Example 4.1, the tree  $\langle\langle \text{Hear\_alarm}(x) \rangle\rangle$  capturing the distribution semantics of  $\text{Hear\_alarm}(x)$  is (partially) depicted as follows. Note the presence

of clauses labelling the clause-nodes.



## 5. WEIGHTED LOGIC PROGRAMMING

Weighted logic programming (WLP) is an abstract framework for describing dynamic programming algorithms in a number of fields, such as natural language processing [EB20] and computational biology [DEKM98]. WLP generalises logic programming by assigning a weight/score to each clause. Given a goal  $A$ , the central task then becomes computing the weight of  $A$ , based on the weights of the clauses appearing in the proof tree of  $A$ .

As WLP and PLP are seemingly close formalisms, it is natural to ask to what extent our coalgebraic perspective applies to WLP. In this section we develop such perspective. By analogy with the development of coalgebraic PLP, we first discuss the ground case (Subsection 5.1), and then move to the general case (Subsection 5.2).


We begin by recalling the basics of WLP, referring to [EB20] for further details. We fix a semiring  $\mathbb{K} = \langle K, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ , where  $\mathbf{0}$  is the  $\oplus$ -unit, and  $\mathbf{1}$  is the  $\otimes$ -unit. A weighted logic program  $\mathbb{W}$  based on a logic program  $\mathbb{L}$  assigns to each clause  $\mathcal{C}$  in  $\mathbb{L}$  a weight label  $c \in K$ . In other words, a weighted logic program  $\mathbb{W}$  consists of finitely many clauses of the form:

$$c :: H \leftarrow B_1, \dots, B_k$$

where  $H, B_1, \dots, B_k$  are atoms and  $c \in K$ .<sup>2</sup>

**Example 5.1.** We use as running example the following weighted logic program  $\mathbb{P}^{sp}$ , from [CSS10].  $\mathbb{P}^{sp}$  is based on the semiring  $\mathbb{K}^{sp} = \langle \mathbb{R}_{\geq 0} \cup \{+\infty\}, \min, +, +\infty, 0 \rangle$ . Its language consists of constants  $a, b, c, d$ , a 1-ary predicate  $\text{initial}(-)$ , and 2-ary predicates  $\text{edge}(-, -)$  and  $\text{reachable}(-, -)$ .

<sup>2</sup>It is typical in the literature (e.g. [CSS10]) to only assign weights to facts, i.e. clauses with empty bodies, and leave the other clauses unlabelled. However, we can assume without loss of generality that all clauses are assigned a weight, as the weight semantics (5.1) implicitly assigns weight 1 to unlabelled clause.



$\mathbb{P}^{sp}$  describes the single-source directed graph on the right. It expresses that a state  $x$  is reachable if either  $x$  itself is the initial state, or there is an edge from a reachable state to  $x$ . For simplicity, we will sometimes refer to the predicates appearing in  $\mathbb{P}^{sp}$  by their abbreviations, such as  $\text{init}(-)$ ,  $\text{reach}(-, -)$ .

As for the semantics, the central task for WLP is computing the *weight*  $\omega^{\mathbb{W}}(A)$  (or simply  $\omega(A)$ ) of an atom  $A$  in  $\mathbb{W}$ . This is an element of  $\mathbb{K}$ , defined as follows:

$$\omega^{\mathbb{W}}(A) = \bigoplus_{\substack{c::A \leftarrow B_1, \dots, B_k \text{ is substitution} \\ \text{instance of some clause in } \mathbb{W}}} \left( c \otimes \bigotimes_{B_i} \omega^{\mathbb{W}}(B_i) \right) \quad (5.1)$$

We will refer to (5.1) as the *weight semantics* for WLP. Note that, when the body of a clause is empty,  $\bigotimes_{B_i} \omega^{\mathbb{W}}(B_i)$  is an empty product and yields the unit  $\mathbf{1}$ . Also, the weight function  $\omega$  is not necessarily well-defined for every goal and program. For example, consider the program consisting of a single clause  $2 :: A \leftarrow A$  with the semiring in Example 5.1. Then (5.1) does not assign a value to  $A$ , because the recursive calculation does not terminate.

**Example 5.2.** In the program  $\mathbb{P}^{sp}$  from Example 5.1, the goal  $\text{reachable}(d)$  has weight  $\omega(\text{reachable}(d)) = 19$ , and it is exactly the weight of the shortest path from  $a$  to  $d$ .

**Remark 5.3.** One might wonder whether PLP in Section 3 and 4 is a special case of WLP. The answer is negative, for two reasons. First, the underlying algebra structure for PLP is  $\langle [0, 1], \vee, \times, 0, 1 \rangle$ , which is not a semiring (distributivity fails). Second, the distribution semantics in (2.1) is not a special case of the weight semantics in (5.1).

**5.1. Coalgebraic Semantics of Ground WLP.** We will present the coalgebraic semantics for ground WLP analogous to that for PLP in Section 3: we first give a coalgebraic presentation of WLP; then we induce a derivation semantics for it. The computation of the weight semantics from the derivation semantics is in Appendix C.

A ground WLP program  $\mathbb{W}$  can be represented as a coalgebra of the type  $\mathcal{MP}_f: \mathbf{Sets} \rightarrow \mathbf{Sets}$ , where  $\mathcal{M}$  is the multiset functor corresponding to the underlying semiring  $\mathbb{K} = \langle K, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ . On objects,  $\mathcal{M}: \mathbf{Sets} \rightarrow \mathbf{Sets}$  maps a set  $A$  to  $\{\varphi: A \rightarrow \mathbb{K} \mid \text{supp}(\varphi) \text{ is finite}\}$ , where  $\text{supp}(\varphi) = \{a \in A \mid \varphi(a) \neq \mathbf{0}\}$ . For each  $a \in A$ , we call  $\varphi(a)$  the weight of  $a$  (under  $\varphi$ ). We adopt the standard notation  $\sum_{i=1}^k c_i a_i$  for the function  $\varphi \in \mathcal{M}(A)$  with support  $\text{supp}(\varphi) = \{a_1, \dots, a_k\}$  and weights  $\varphi(a_i) = c_i$ . We also use  $\emptyset$  for the multiset with empty support. On morphisms,  $\mathcal{M}(h: A \rightarrow B)$  maps  $\sum_{i=1}^k c_i a_i$  to  $\sum_{i=1}^k c_i h(a_i)$ .

Fix a ground WLP program  $\mathbb{W}$  based on a semiring  $\mathbb{K}$  and a set of atoms  $\text{At}$ . The definition of  $\mathbb{W}$  can be encoded as an  $\mathcal{MP}_f$ -coalgebra  $w: \text{At} \rightarrow \mathcal{MP}_f(\text{At})$  as follows: for

each  $A \in \text{At}$ ,


$$w(A): \quad \mathcal{P}_f(\text{At}) \quad \rightarrow \quad K$$

$$\{B_1, \dots, B_n\} \mapsto \begin{cases} c & \text{if } c :: A \leftarrow B_1, \dots, B_n \text{ is a clause in } \mathbb{W} \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

**Example 5.4.** Consider the following ground program  $\mathbb{P}^{gsp}$  obtained as a simple ground version of  $\mathbb{P}^{sp}$ . The underlying semiring is still  $\mathbb{K}^{sp} = \langle \mathbb{R}_{\geq 0} \cup \{+\infty\}, \min, +, +\infty, 0 \rangle$ .

$$\begin{array}{l|l} 0 :: \text{initial}(a) & \leftarrow \\ 4 :: \text{edge}(a, c) & \leftarrow \\ 20 :: \text{edge}(a, d) & \leftarrow \\ 9 :: \text{edge}(c, a) & \leftarrow \\ 15 :: \text{edge}(c, d) & \leftarrow \\ 16 :: \text{edge}(d, c) & \leftarrow \\ 0 :: \text{reachable}(a) & \leftarrow \text{initial}(a) \\ 0 :: \text{reachable}(c) & \leftarrow \text{initial}(c) \end{array} \quad \left| \quad \begin{array}{l} 0 :: \text{reachable}(d) \leftarrow \text{initial}(d) \\ 0 :: \text{reachable}(a) \leftarrow \text{reachable}(c), \text{edge}(c, a) \\ 0 :: \text{reachable}(a) \leftarrow \text{reachable}(d), \text{edge}(d, a) \\ 0 :: \text{reachable}(c) \leftarrow \text{reachable}(a), \text{edge}(a, c) \\ 0 :: \text{reachable}(c) \leftarrow \text{reachable}(d), \text{edge}(d, c) \\ 0 :: \text{reachable}(d) \leftarrow \text{reachable}(a), \text{edge}(a, d) \\ 0 :: \text{reachable}(d) \leftarrow \text{reachable}(c), \text{edge}(c, d) \end{array} \right.$$

The program describes the (single-sourced, directed) weighted graph (5.2) on the right:



The (finite) set of atoms  $\text{At}_{gsp}$  is

$$\{\text{initial}(a), \dots, \text{initial}(c), \text{edge}(a, c), \dots, \text{edge}(d, c), \text{reachable}(a), \dots, \text{reachable}(c)\}.$$

Some weights of the corresponding coalgebra  $p^{gsp}: \text{At}_{gsp} \rightarrow \mathcal{MP}_f(\text{At}_{gsp})$  are presented below:


$$\begin{aligned} p^{gsp}(\text{initial}(a)) &= 0\{ \} & p^{gsp}(\text{edge}(a, d)) &= 20\{ \} & p^{gsp}(\text{edge}(d, a)) &= \emptyset \\ p^{gsp}(\text{reachable}(a)) &= 0\{\text{initial}(a)\} + 0\{\text{reachable}(c), \text{edge}(c, a)\} + 0\{\text{reachable}(d), \text{edge}(d, a)\} \end{aligned}$$

Note that the 0 in the above equations is the  $\otimes$ -unit  $\mathbf{1}$  in  $\mathbb{K}^{sp}$ , and should not be confused with the  $\oplus$ -unit  $\mathbf{0}$ .

Just as in the case of PLP, we may represent operationally the recursive calculation of the weight  $\omega(A)$  of a goal  $A$  as a certain kind of tree, which we call *weighted derivation tree*. These are weighted version of the and-or trees for pure logic programming [GC94]. They are formally defined analogously to the stochastic derivation trees in Definition 3.3, except that the probability labels are replaced by weight labels. We illustrate the concept with an example.

**Example 5.5.** In the context of Example 5.4, the weighted derivation tree for  $\text{reachable}(a)$  is partially depicted below. Note the different meaning for an atom-node and a clause-node to have no child. For instance, the clause-node following  $\text{init}(a)$  has no child because ' $0 :: \text{initial}(a) \leftarrow$ ' is a clause in  $\mathbb{P}^{gsp}$ , while the atom-node  $\text{edge}(d, a)$  has no child because

there is no clause in  $\mathbb{P}^{gsp}$  whose head is  $\text{edge}(d, a)$ .



By replacing  $\mathcal{M}_{pr}$  by  $\mathcal{M}$  in Construction 3.5, one may construct the terminal sequence for the functor  $\text{At} \times \mathcal{MP}_f(-)$ , which by accessibility [Wor99] converges to a limit  $X_\gamma$ , yielding the final  $\text{At} \times \mathcal{MP}_f$ -coalgebra  $X_\gamma \cong \text{At} \times \mathcal{MP}_f(X_\gamma)$  and a unique coalgebra morphism  $\llbracket - \rrbracket_w$  as below:

$$\begin{array}{ccc} \text{At} & \xrightarrow{\llbracket - \rrbracket_w} & X_\gamma \\ \downarrow \langle id, w \rangle & \searrow id \times \mathcal{MP}_f(\llbracket - \rrbracket_w) & \downarrow \cong \\ \text{At} \times \mathcal{MP}_f(\text{At}) & \longrightarrow & \text{At} \times \mathcal{MP}_f(X_\gamma). \end{array}$$

By construction, one may readily verify that weighted derivation trees are elements of  $X_\gamma$ , and  $\llbracket - \rrbracket_w$  maps an atom  $A$  to its weighted derivation tree, computed according to the program  $\mathbb{W}$ . Thus we can define the derivation semantics for WLP as the map  $\llbracket - \rrbracket_w$ . Note the weight  $\omega^w(A)$  of a goal is effectively computable from its semantics  $\llbracket A \rrbracket_w$  — see Appendix C for details.

**5.2. Coalgebraic Semantics for Arbitrary WLP Programs.** In this subsection we generalise our approach to arbitrary weighted logic programs, possibly including variables and functions. We shall assume that the semiring  $\mathbb{K}$  underlying the WLP programs is countably  $\oplus$ -complete, namely  $(\bigoplus_{c \in C} c) \in K$  and  $d \otimes (\bigoplus_{c \in C} c) = \bigoplus_{c \in C} (d \otimes c)$ , for arbitrary countable subset  $C \subseteq K$ . As we will see, the reason for such assumption is that there may exist countably many substitutions that match the head of some clause to the goal, generating countably many different instances of the body. The weight semantics requires adding up all these countably many weights, whence their sum should be an element of  $\mathbb{K}$ .


The construction leading to a coalgebraically defined derivation semantics  $\llbracket - \rrbracket$  for WLP is fairly similar to the one for PLP (Subsection 4.2), thus we confine ourselves to highlighting the main steps and emphasise the differences between the two approaches.

The derivation semantics  $\llbracket - \rrbracket$  will associate a goal with its resolution by unification in a program. This resolution will be represented as a tree, which we call *weighted saturated derivation tree*. Formally, these are defined analogously to stochastic saturated derivation trees (Definition 4.7), with the difference that there is no need of clause-nodes (for reasons detailed in Remark 5.7 below), and labels on edges departing from substitution-nodes are now weights rather than probabilities.

Before providing a coalgebraic modelling of the semantics, we establish some preliminary intuition via our leading example.

**Example 5.6.** Recall program  $\mathbb{P}^{sp}$  from Example 5.1. We show below part of the weighted saturated derivation tree for the goal  $\text{reachable}(x)$  (in context 1), where we write  $\emptyset$  for the empty substitution. Compared with the ground case (see (5.3) in Example 5.5), the tree in (5.4) has an extra layer of ‘substitution-nodes’, labelled with the substitutions applied to the goal during resolution by unification. For instance, the node labelled with  $x \mapsto \mathbf{b}$  has one child for each substitution instance of the body  $\{B_1, \dots, B_k\}\tau$  satisfying that  $c :: H \leftarrow B_1, \dots, B_k$

is a clause in  $\mathbb{P}^{sp}$  and the head  $H$  term-matches the goal  $\text{reach}(x)[x \mapsto b] = \text{reach}(b)$  via substitution  $\tau$ .



As mentioned, with respect to stochastic saturated derivation trees (see e.g. (4.1)), the semantics represented in (5.4) does not require the presence of clause-nodes (which were labelled with  $\bullet$  in (4.1)).

Similarly to the case of PLP, we will describe resolution by unification in two steps: first, one considers all substitutions  $\theta$  of the goal  $A$ . Then, each substitution instance  $A\theta$  of the goal is compared by term-matching to the heads of clauses in the program  $\mathbb{W}$ . The function  $u$  performing term-matching of  $A\theta$  in  $\mathbb{W}$  associates a set  $\{C_1, \dots, C_m\}$  of atoms with an element of  $\mathbb{K}$ , computed as follows:

$$u(A\theta): \{C_1, \dots, C_m\} \mapsto \bigoplus_{\substack{c :: H \leftarrow B_1, \dots, B_k \text{ is a clause of } \mathbb{W} \\ \text{such that } \exists \tau \text{ with } H\tau = A\theta \\ \text{and } \{B_1, \dots, B_k\}\tau = \{C_1, \dots, C_m\}}} c \quad (5.5)$$

Intuitively, for each clause  $c :: H \leftarrow B_1, \dots, B_k$  in  $\mathbb{W}$  and substitution  $\tau$  that matches the head  $H$  with  $A\theta$ ,  $u(A\theta)$  assigns weight  $c$  to  $\{B_1, \dots, B_k\}\tau$ . There are three observations to make here. First, it is possible that there is a different clause  $c' :: H' \leftarrow B'_1, \dots, B'_{k'}$  and substitution  $\tau'$  such that  $H'\tau' = A\theta$ , and  $\{B'_1, \dots, B'_{k'}\}\tau' = \{B_1, \dots, B_k\}\tau = \{C_1, \dots, C_m\}$ . In this case, the weight  $c'$  is also added to the value of  $u(A\theta)(\{C_1, \dots, C_m\})$ . Second, there may exist a different substitution  $\sigma$  such that  $H\tau = H\sigma$  and  $\{B_1, \dots, B_k\}\tau = \{B_1, \dots, B_k\}\sigma$ . In this case, no extra weight  $c$  is added to  $u(A\theta)(\{B_1, \dots, B_k\}\tau)$ . Third, the support  $\text{supp}(u(A\theta))$  of the function might be countably infinite. This is due to the possible existence of a clause  $H \leftarrow B_1, \dots, B_k$  and countably many substitutions  $\tau_1, \tau_2, \dots$  such that  $H\tau_i = A\theta$  while  $\{B_1, \dots, B_k\}\tau_i \neq \{B_1, \dots, B_k\}\tau_j$ , for all distinct  $i, j = 1, 2, \dots$ .

For example, the application of  $u$  to  $\text{reach}(b)$  in Example 5.6 can be visualised in tree form as the first two layers of the subtree starting at  $x \mapsto b$ .

Towards a categorical treatment of resolution by unification, we need to introduce the *countable* (because of the third observation above) multiset functor  $\mathcal{M}_c: \mathbf{Sets} \rightarrow \mathbf{Sets}$  over a  $\oplus$ -complete semiring  $\mathbb{K}$ . On objects,  $\mathcal{M}_c(X) = \{\varphi: X \rightarrow \mathbb{K} \mid \text{supp}(\varphi) \text{ is at most countable}\}$ . On morphisms,  $\mathcal{M}_c(h): \mathcal{M}_c(X) \rightarrow \mathcal{M}_c(Y)$  maps  $\varphi$  to  $\lambda(y \in Y). \bigoplus_{x \in h^{-1}(y)} \varphi(x)$ . Now, lifting functors from  $\mathbf{Sets}$  to  $\mathbf{Sets}^{\mathbf{L}_{\Sigma}^{\text{op}}}$  as in (4.4), the term-matching part  $u$  of the resolution can be expressed as a  $\widehat{\mathcal{M}}_c \widehat{\mathcal{P}}_f$ -coalgebra in  $\mathbf{Sets}^{\mathbf{L}_{\Sigma}^{\text{op}}}$ :

$$u: \mathcal{U}\text{At} \rightarrow \widehat{\mathcal{M}}_c \widehat{\mathcal{P}}_f \mathcal{U}\text{At} \quad (5.6)$$

**Remark 5.7.** Note the coalgebra type of term-matching for WLP ( $\widehat{\mathcal{M}}_c \widehat{\mathcal{P}}_f$ ) is simpler than the one for PLP ( $\widehat{\mathcal{M}}_{pr} \widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f$ ), as in (4.4). In fact, in the distribution semantics of PLP, the use of a clause in a proof should be counted at most once. This is taken care by the extra layer  $\widehat{\mathcal{M}}_{pr}$  which allows to identify different substitution instances of the same clause in a single

tree. Instead, in the weight semantics every use of a clause counts towards calculating the total weight (*cf.* (5.1)). Thus it is unnecessary to keep track of whether clauses are used more than once.

This difference is reflected in the tree representation of the two semantics, compare e.g. (5.4) with (4.1). In the weighted saturated trees, there is no need for clause-nodes, which are instead present in stochastic saturated trees.

Analogously to the PLP case, once term matching is modelled as a coalgebra, we can then exploit the unit  $\eta$  of the adjunction  $\mathcal{U} \dashv \mathcal{K}$  (see (4.3)) to model the retrieval of all the substitutions of the goals. This will yield a  $\mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}\mathcal{A}t$ -coalgebra in  $\mathbf{Sets}^{\mathbf{L}_\Sigma^{\text{op}}}$ , capturing resolution by unification:

$$\mathbf{u}^\sharp := \mathbf{A}t \xrightarrow{\eta_{\mathbf{A}t}} \mathcal{K}\mathcal{U}\mathbf{A}t \xrightarrow{\mathcal{K}\mathbf{u}} \mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}\mathbf{A}t \quad (5.7)$$

Spelling out the definitions, given some  $A \in \mathbf{A}t(n)$ ,

$$\mathbf{u}_n^\sharp : A \mapsto \langle \mathbf{u}_m(A\theta) \rangle_{\theta \in \mathbf{L}_\Sigma^{\text{op}}[n,m]}$$

Finally, we can formulate the derivation semantics for general WLP via the cofree coalgebra. First, one may generate the terminal sequence for the functor  $\mathbf{A}t \times \mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}$ , following a similar procedure as Construction 4.5. Since  $\mathcal{M}_c$  is accessible, the terminal sequence converges, say at  $Z_\gamma$ . Then the derivation semantics  $\llbracket - \rrbracket_{\mathbf{u}^\sharp}$  is defined by the universal mapping property of the final coalgebra:

$$\begin{array}{ccc} \mathbf{A}t & \xrightarrow{\llbracket - \rrbracket_{\mathbf{u}^\sharp}} & Z_\gamma \\ \downarrow \langle id, \mathbf{u}^\sharp \rangle & & \downarrow \cong \\ \mathbf{A}t \times \mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}\mathbf{A}t & \xrightarrow{id \times \mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}\llbracket - \rrbracket_{\mathbf{u}^\sharp}} & \mathbf{A}t \times \mathcal{K}\widehat{\mathcal{M}}_c\widehat{\mathcal{P}}_f\mathcal{U}Z_\gamma \end{array} \quad (5.8)$$

As expected, weighted saturated derivation trees can be regarded as elements of the final coalgebra  $Z_\gamma$ , and  $\llbracket - \rrbracket_{\mathbf{u}^\sharp}$  maps an atom  $A$  to its weighted saturated derivation tree.

## 6. CONCLUSION

This work proposed a coalgebraic semantics for probabilistic logic programming and weighted logic programming, as an extension of the framework in [KMP10, BZ15] for pure logic programming. Our approach consists of the following three steps. First we represent the programs as coalgebras of appropriate type. Next we define the derivation semantics in terms of certain derivation trees, which can be recovered as objects in the corresponding final coalgebra. Finally, we explain how to retrieve the semantics that is most commonly found in the literature (distribution semantics for PLP and weight semantics for WLP) from our coalgebraic semantics. While the first two steps simply generalise the approach in [BZ15], the last step has no counterpart in pure logic programming, and it requires extra categorical constructions (distribution trees) and computations (Appendix A, B, C).

In extending our coalgebraic semantics from ground to arbitrary PLP programs, we adopted the ‘saturated’ approach of [BZ15]. The ‘lax naturality’ approach of [KP11] offers an alternative route (see [KP18a] for a detailed comparison), based on the observation that, even though first-order logic programs fail to be natural transformations, they are lax natural in a suitable order enriched category. Developing a coalgebraic lax semantics for PLP and WLP would yield derivation trees representing computations by term-matching

rather than general unification. The extra information given by lax naturality could be used for algorithmic purposes, similarly to what is done in [KSH14, KPS13] for pure logic programming. We leave the development of lax semantics for PLP and WLP as future work.

As follow-up work, we plan to investigate in two main directions. First, starting from the observation that Bayesian networks (over binary variables) can be seen as PLP programs, we want to understand Bayesian reasoning within our coalgebraic framework. Recent works on a categorical semantics for Bayesian probability (in particular [JZ19, CJ19]) provide a starting point for this research. A major challenge is that standard translations of Bayesian networks into PLP (such as the one in [Poo08]) require to incorporate negation in the logical syntax. Thus a preliminary step will be to model coalgebraically PLP with negation, which we plan to do following the credal set approach outlined in [Coz19].

A second research thread is investigating the semantics of coinductive logic programs in the quantitative setting of PLP and WLP. This is an essentially unexplored perspective, which we will tackle building on research on coinduction in the coalgebraic semantics of standard logic programming [KL18, GBM<sup>+</sup>07, BKL19].

## REFERENCES

- [BKL19] Henning Basold, Ekaterina Komendantskaya, and Yue Li. Coinduction in uniform: Foundations for corecursive proof search with horn clauses. *Lecture Notes in Computer Science*, pages 783–813. Springer, 2019.
- [BSdV04] Falk Bartels, Ana Sokolova, and Erik P. de Vink. A hierarchy of probabilistic system types. *Theor. Comput. Sci.*, 327(1-2):3–22, 2004.
- [BZ13] Filippo Bonchi and Fabio Zanasi. Saturated semantics for coalgebraic logic programming. In *Algebra and Coalgebra in Computer Science - 5th International Conference, CALCO 2013, Warsaw, Poland, September 3-6, 2013. Proceedings*, pages 80–94, 2013.
- [BZ15] Filippo Bonchi and Fabio Zanasi. Bialgebraic semantics for logic programming. *Logical Methods in Computer Science*, 11(1), 2015.
- [CJ19] Kenta Cho and Bart Jacobs. Disintegration and bayesian inversion via string diagrams. *Mathematical Structures in Computer Science*, 29(7):938–971, 2019.
- [Coz19] Fabio Cozman. The joy of probabilistic answer set programming. In Jasper De Bock, Cassio P. de Campos, Gert de Cooman, Erik Quaeghebeur, and Gregory Wheeler, editors, *Proceedings of the Eleventh International Symposium on Imprecise Probabilities: Theories and Applications*, volume 103 of *Proceedings of Machine Learning Research*, pages 91–101, Thagaste, Ghent, Belgium, 03–06 Jul 2019. PMLR.
- [CSS10] Shay Cohen, Robert Simmons, and Noah Smith. Products of weighted logic programs. *Computing Research Repository - CORR*, 11, 06 2010.
- [Dan92] Eugene Dantsin. Probabilistic logic programs and their semantics. In A. Voronkov, editor, *Logic Programming*, pages 152–164, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.
- [DDGK16] Fredrik Dahlqvist, Vincent Danos, Ilias Garnier, and Ohad Kammar. Bayesian inversion by  $\omega$ -complete cone duality. In *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, pages 1:1–1:15, 2016.
- [DEKM98] Richard Durbin, Sean R. Eddy, Anders Krogh, and Graeme J. Mitchison. *Biological sequence analysis: Probabilistic models of proteins and nucleic acids*. Cambridge University Press, 1998.
- [DRKT07a] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: A probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 2468–2473, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- [DRKT07b] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: A probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 2468–2473, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.



- [EB20] Jason Eisner and John Blatz. Program transformations for optimization of parsing algorithms and other weighted logic programs. 02 2020.
- [GBM<sup>+</sup>07] Gopal Gupta, Ajay Bansal, Richard Min, Luke Simon, and Ajay Mallya. Coinductive logic programming and its applications. In Véronica Dahl and Ilkka Niemelä, editors, *Logic Programming*, pages 27–44, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- [GC94] Gopal Gupta and Vítor Santos Costa. Optimal implementation of and-or parallel prolog. *Future Generation Computer Systems*, 10(1):71 – 92, 1994. PARLE '92.
- [GZ19] Tao Gu and Fabio Zanasi. A coalgebraic perspective on probabilistic logic programming. In *8th Conference on Algebra and Coalgebra in Computer Science, CALCO 2019, June 3-6, 2019, London, United Kingdom*, pages 10:1–10:21, 2019.
- [JKZ19] Bart Jacobs, Aleks Kissinger, and Fabio Zanasi. Causal inference by string diagram surgery. In *Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Proceedings*, 2019.
- [JZ19] Bart Jacobs and Fabio Zanasi. The logical essentials of bayesian reasoning. In Joost-Peter Katoen Gilles Barthe and Alexandra Silva, editors, *Probabilistic Programming*. Cambridge University Press, Cambridge, 2019.
- [KL17] Ekaterina Komendantskaya and Yue Li. Productive corecursion in logic programming. *TPLP*, 17(5-6):906–923, 2017.
- [KL18] Ekaterina Komendantskaya and Yue Li. Towards coinductive theory exploration in horn clause logic: Position paper. In *Proceedings 5th Workshop on Horn Clauses for Verification and Synthesis, HCVS 2018, Oxford, UK, 13th July 2018.*, pages 27–33, 2018.
- [KMP10] Ekaterina Komendantskaya, Guy McCusker, and John Power. Coalgebraic semantics for parallel derivation strategies in logic programming. In *Algebraic Methodology and Software Technology - 13th International Conference, AMAST 2010, Lac-Beauport, QC, Canada, June 23-25, 2010. Revised Selected Papers*, pages 111–127, 2010.
- [KP11] Ekaterina Komendantskaya and John Power. Coalgebraic semantics for derivations in logic programming. In *Algebra and Coalgebra in Computer Science - 4th International Conference, CALCO 2011, Winchester, UK, August 30 - September 2, 2011. Proceedings*, pages 268–282, 2011.
- [KP18a] Ekaterina Komendantskaya and John Power. Logic programming: laxness and saturation. *Journal of logical and algebraic methods in programming*, 101:1–21, 2018.
- [KP18b] Ekaterina Komendantskaya and John Power. Logic programming: Laxness and saturation. *J. Log. Algebr. Meth. Program.*, 101:1–21, 2018.
- [KPS13] Ekaterina Komendantskaya, John Power, and Martin Schmidt. Coalgebraic logic programming: from semantics to implementation. *CoRR*, abs/1312.6568, 2013.
- [KPS16] Ekaterina Komendantskaya, John Power, and Martin Schmidt. Coalgebraic logic programming: from semantics to implementation. *J. Log. Comput.*, 26(2):745–783, 2016.
- [KSH14] Ekaterina Komendantskaya, Martin Schmidt, and Jónathan Heras. Exploiting parallelism in coalgebraic logic programming. *Electron. Notes Theor. Comput. Sci.*, 303:121–148, 2014.
- [Llo87] John W. Lloyd. *Foundations of Logic Programming, 2nd Edition*. Springer, 1987.
- [Mac71] Saunders MacLane. *Categories for the Working Mathematician*. Springer-Verlag, 1971. Graduate Texts in Mathematics, Vol. 5.
- [MH12] Søren Mørk and Ian Holmes. Evaluating bacterial gene-finding hmm structures as probabilistic logic programs. *Bioinformatics*, 28(5):636–642, 2012.
- [NS92] Raymond Ng and V.S. Subrahmanian. Probabilistic logic programming. *Information and Computation*, 101(2):150 – 201, 1992.
- [Poo08] David Poole. The independent choice logic and beyond. In Luc De Raedt, Paolo Frasconi, Kristian Kersting, and Stephen Muggleton, editors, *Probabilistic Inductive Logic Programming - Theory and Applications*, volume 4911 of *Lecture Notes in Computer Science*, pages 222–243. Springer, 2008.
- [RS14] Fabrizio Riguzzi and Terrance Swift. Probabilistic logic programming under the distribution semantics. 2014.

- [Sat95] Taisuke Sato. A statistical learning method for logic programs with distribution semantics. In *Logic Programming, Proceedings of the Twelfth International Conference on Logic Programming, Tokyo, Japan, June 13-16, 1995*, pages 715–729, 1995.
- [Sil10] Alexandra Silva. Kleene coalgebra. Phd thesis, CWI, Amsterdam, The Netherlands, 2010.
- [TBF05] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT Press, Cambridge, Mass., 2005.
- [Wor99] James Worrell. Terminal sequences for accessible endofunctors. *Electronic Notes in Theoretical Computer Science*, 19:24 – 38, 1999. CMCS’99, Coalgebraic Methods in Computer Science.
- [Zes17] Riccardo Zese. *Probabilistic Semantic Web: Reasoning and Learning*. IOS Press, Amsterdam, The Netherlands, The Netherlands, 2017.

## APPENDIX A. COMPUTABILITY OF THE DISTRIBUTION SEMANTICS (GROUND CASE)

**Computing with distribution trees.** As a justification for our tree representation of the distribution semantics, we claimed that the probability  $\Pr_{\mathbb{P}}(A)$  associated with a goal (see (2.1)) can be straightforwardly computed from the corresponding distribution tree  $\langle\langle A \rangle\rangle_p$ . This appendix supplies such an algorithm. Note that this serves just as a proof of concept, without any claim of efficiency compared to pre-existing implementations. In this section we fix a ground PLP program  $\mathbb{P}$  with atoms  $\text{At}$ , a goal  $A \in \text{At}$  and the distribution tree  $\mathcal{T}$  for  $A$  in  $\mathbb{P}$  (Definition 3.8). First, we may assume that the distribution tree  $\mathcal{T}$  is finite. Indeed, in the ground case infinite branches only result from loops, namely multiple appearances of an atom in some path, which can be easily detected. We can always prune the subtrees of  $\mathcal{T}$  rooted by atoms that already appeared at an earlier stage: this does not affect the computation of  $\Pr_{\mathbb{P}}(A)$ , and it makes  $\mathcal{T}$  finite. Next, we introduce the concept of *deterministic* subtree. Basically a deterministic subtree selects one world-node at each stage. Recall that every clause-node in  $\mathcal{T}$  represents a clause in  $|\mathbb{P}|$ , whose head is the label of its atom-grandparent, and body consists of the labels of its atom-children.

**Definition A.1.** A subtree  $\mathcal{S}$  of  $\mathcal{T}$  is *deterministic* if (i) it contains exactly one child (world-node) for each atom-node and all children for other nodes, and (ii) for any distinct atom-nodes  $s, t$  in  $\mathcal{S}$  with the same label,  $s$  and  $t$  have their clause-grandchildren representing the same clauses.

The idea is that  $\mathcal{S}$  describes a computation in which the choice of a possible world (i.e., a sub-program of  $\mathbb{P}$ ) associated to any atom  $B$  appearing during the resolution is uniquely determined. Because of this feature, each deterministic subtree uniquely identifies a set of sub-programs of  $\mathbb{P}$ , and together the deterministic subtrees of  $\mathcal{T}$  form a *partition* over the set of these sub-programs (see Proposition A.3 below).

Since  $\mathcal{T}$  is finite, it is clear that we can always provide an enumeration of its deterministic subtrees. We can now present our algorithm, in two steps. First, Algorithm 1 computes the probability associated with a deterministic subtree. Second, Algorithm 2 computes  $\Pr_{\mathbb{P}}(A)$  by summing up the probabilities found by Algorithm 1 on all the deterministic subtrees of  $\mathcal{T}$  that contains a refutation of  $A$ . Below we write ‘label( $s \rightarrow t$ )’ for the probability value labelling the edge from  $s$  to  $t$ .

---

**Algorithm 1** Compute probability of a deterministic subtree

---

**Input:** A deterministic subtree  $\mathcal{S}$  of  $\mathcal{T}$

**Output:** The probability of  $\mathcal{S}$

```

1: prob_list = []
2: for atom-node  $s$  in  $\mathcal{S}$  do
3:   if  $s$  has child then
4:     prob_list.append(label( $s \rightarrow$  child( $s$ )))
5: if prob_list == [] then
6:   return 0
7: else prob = product of values in prob_list
8:   return prob

```

---

---

**Algorithm 2** Compute probability of a goal
 

---

**Input:** The distribution tree  $\mathcal{T}$  of  $A$  in  $\mathbb{P}$

**Output:** The success probability  $\Pr_{\mathbb{P}}(A)$

```

1: prob_suc = 0
2: for deterministic subtree  $\mathcal{S}$  of  $\mathcal{T}$  do
3:   if  $\mathcal{S}$  proves  $A$  then
4:     prob_suc += Algorithm 1( $\mathcal{S}$ )
5: return prob_suc
  
```

---

The above procedure terminates because  $\mathcal{T}$  is finite and every for-loop is finite. We now focus on the correctness of the algorithm.

**Correctness.** As mentioned, a world-node in a deterministic subtree can be seen as a choice of clauses: one chooses the clauses represented by its clause-children, and discards the clauses represented by its ‘complement’ world. For correctness, we make this precise, via the following definition.

**Definition A.2.** Given a clause  $\mathcal{C}$  in  $\mathbb{P}$ , a deterministic subtree  $\mathcal{S}$  of  $\mathcal{T}$ , a world-node  $t$  and its atom-parent  $s$  in  $\mathcal{S}$ , we say  $t$  *accepts*  $\mathcal{C}$  if  $\text{Head}(\mathcal{C}) = \text{label}(s)$  and there is a clause-child of  $t$  that represents  $\mathcal{C}$ ;  $t$  *rejects*  $\mathcal{C}$  if  $\text{Head}(\mathcal{C}) = \text{label}(s)$  but no clause-child of  $t$  represents  $\mathcal{C}$ . We say  $\mathcal{S}$  *accepts* (*rejects*)  $\mathcal{C}$  if there exists a world-node  $t$  in  $\mathcal{S}$  that accepts (rejects)  $\mathcal{C}$ .

Note that Definition A.1, condition (ii) prevents the existence of world-nodes  $t, t'$  in  $\mathcal{S}$  such that  $t$  accepts  $\mathcal{C}$  and  $t'$  rejects  $\mathcal{C}$ . Thus the notion that  $\mathcal{S}$  accepts (rejects)  $\mathcal{C}$  is well-defined. We denote the set of clauses accepted and rejected by  $\mathcal{S}$  by  $\text{Acc}(\mathcal{S})$  and  $\text{Rej}(\mathcal{S})$ , respectively. Then we can define the set  $\text{SubProg}(\mathcal{S})$  of sub-programs represented by  $\mathcal{S}$  as

$$\text{SubProg}(\mathcal{S}) := \{\mathbb{L} \subseteq |\mathbb{P}| \mid \forall \mathcal{C} \in \text{Acc}(\mathcal{S}), \mathcal{C} \in \mathbb{L}; \forall \mathcal{C}' \in \text{Rej}(\mathcal{S}), \mathcal{C}' \notin \mathbb{L}\} \quad (\text{A.1})$$

We will prove the correctness of the algorithm through the following basic observations on the connection between deterministic subtrees and the sub-programs they represent:

**Proposition A.3.** *Suppose  $\mathcal{S}$  is a deterministic subtree of the distribution tree  $\mathcal{T}$  of  $A$ .*

- (1)  $\{\text{SubProg}(\mathcal{S}) \mid \mathcal{S} \text{ is deterministic subtree of } \mathcal{T}\}$  forms a partition of  $\mathcal{P}(\mathbb{P})$ .
- (2) Either  $\mathbb{L} \vdash A$  for all  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$  or  $\mathbb{L} \not\vdash A$  for all  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$ .
- (3)  $\sum_{\mathbb{L} \in \text{SubProg}(\mathcal{S})} \Pr_{\mathbb{P}}(\mathbb{L}) = \prod_{r_i \in \mathcal{S}} r_i$ , where the  $r_i$ s are all the probability labels appearing in  $\mathcal{S}$  (on the atom-node  $\rightarrow$  world-node edges).

*Proof.*

- (1) Given any two distinct deterministic subtrees, there is an atom-node  $s$  such that the subtrees include distinct world-child of  $s$ . So by (A.1) the sub-programs they represent do not share at least one clause. Moreover, given a sub-program  $\mathbb{L}$ , one can always identify a deterministic subtree  $\mathcal{S}$  such that  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$ , as follows: given the  $A$ -labelled root of  $\mathcal{T}$ , select the world-child  $w$  of  $A$  representing the (possibly empty) set  $X$  of all clauses in  $\mathbb{L}$  whose head is  $A$ ; then select the children (if any) of  $w$ , and repeat the procedure.
- (2) Note that a sub-program  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$  proves the goal  $A$  iff  $\mathcal{S}$  contains a successful refutation of  $A$ , and the latter property is independent of the choice of  $\mathbb{L}$ .

(3) We refer to  $\prod_{r_i \in \mathcal{S}} r_i$  as the probability of the deterministic subtree  $\mathcal{S}$ . For each sub-program  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$ , its probability can be written as

$$\mu_{\mathbb{P}}(\mathbb{L}) = \left( \prod_{\mathcal{C} \in \text{Acc}(\mathcal{S})} \text{Label}(\mathcal{C}) \right) \cdot \left( \prod_{\mathcal{C}' \in \text{Rej}(\mathcal{S})} (1 - \text{Label}(\mathcal{C}')) \right) \cdot \mu_{\mathbb{P} \setminus (\text{Acc} \cup \text{Rej})}(\mathbb{L} \setminus \text{Acc}(\mathcal{S})) \quad (\text{A.2})$$

Note that  $\text{SubProg}(\mathcal{S})$  can also be written as  $\{X \cup \text{Acc}(\mathcal{S}) \mid X \subseteq \mathbb{P} \setminus (\text{Acc}(\mathcal{S}) \cup \text{Rej}(\mathcal{S}))\}$ , so

$$\sum_{\mathbb{L} \in \text{SubProg}(\mathcal{S})} \mu_{\mathbb{P} \setminus (\text{Acc} \cup \text{Rej})}(\mathbb{L} \setminus \text{Acc}(\mathcal{S})) = 1. \quad (\text{A.3})$$

Applying equation (A.3) to the sum of (A.2) over all  $\mathbb{L} \in \text{SubProg}(\mathcal{S})$ , we get

$$\sum_{\mathbb{L} \in \text{SubProg}(\mathcal{S})} \mu_{\mathbb{P}}(\mathbb{L}) = \prod_{\mathcal{C} \in \text{Acc}(\mathcal{S})} \text{Label}(\mathcal{C}) \cdot \prod_{\mathcal{C}' \in \text{Rej}(\mathcal{S})} (1 - \text{Label}(\mathcal{C}')) \quad (\text{A.4})$$

For each world-node  $t$  and its atom-parent  $s$ , we can use the terminology in Definition A.2, and express  $\text{label}(s \rightarrow t)$  (see Definition 3.8) as

$$\text{label}(s \rightarrow t) = \prod_{t \text{ accepts } \mathcal{C}} \text{Label}(\mathcal{C}) \cdot \prod_{t \text{ rejects } \mathcal{C}'} (1 - \text{Label}(\mathcal{C}')). \quad (\text{A.5})$$

Applying (A.5) to the whole deterministic subtree  $\mathcal{S}$ , we obtain

$$\begin{aligned} \sum_{\mathbb{L} \in \text{SubProg}(\mathcal{S})} \mu_{\mathbb{P}}(\mathbb{L}) &\stackrel{(\text{A.4})}{=} \prod_{\mathcal{C} \in \text{Acc}(\mathcal{S})} \text{Label}(\mathcal{C}) \cdot \prod_{\mathcal{C}' \in \text{Rej}(\mathcal{S})} (1 - \text{Label}(\mathcal{C}')) \\ &\stackrel{\text{Def. A.2}}{=} \prod_{(\text{world-node } t \text{ in } \mathcal{S})} \left( \prod_{t \text{ accepts } \mathcal{C}} \text{Label}(\mathcal{C}) \cdot \prod_{t \text{ rejects } \mathcal{C}'} (1 - \text{Label}(\mathcal{C}')) \right) \\ &\stackrel{(\text{A.5})}{=} \prod_{r_i \in \mathcal{S}} r_i \end{aligned}$$

If we say two world-nodes  $t$  and  $t'$  are equivalent if their clause-children represent exactly the same clauses in  $\mathbb{P}$ , then the  $\prod_{(\text{world-node } t \text{ in } \mathcal{S})}$  in the above calculation visits every world-node exactly once modulo equivalence.  $\square$

We can now formulate the success probability of  $A$  as follows

$$\begin{aligned} \text{Pr}_{\mathbb{P}}(A) &= \sum_{|\mathbb{P}| \supseteq \mathbb{L} \vdash A} \mu_{\mathbb{P}}(\mathbb{L}) \stackrel{(\text{Prop. A.3, 1\&2})}{=} \sum_{S \vdash A} \sum_{\mathbb{L} \in \text{SubProg}(\mathcal{S})} \mu_{\mathbb{P}}(\mathbb{L}) \\ &\stackrel{(\text{A.4})}{=} \sum_{S \vdash A} \left( \prod_{\mathcal{C} \in \text{Acc}(\mathcal{S})} \text{Label}(\mathcal{C}) \cdot \prod_{\mathcal{C}' \in \text{Rej}(\mathcal{S})} (1 - \text{Label}(\mathcal{C}')) \right) \stackrel{(\text{Prop. A.3, 3})}{=} \sum_{S \vdash A} \prod_{r_i \in \mathcal{S}} r_i \end{aligned}$$

In words, this is exactly Algorithm 2: we sum up the probabilities of all deterministic subtrees  $\mathcal{S}$  of the distribution tree  $\mathcal{T}$  which contain a proof of  $A$ .

## APPENDIX B. COMPUTABILITY OF THE DISTRIBUTION SEMANTICS (GENERAL CASE)

Computability of the distribution semantics for arbitrary PLP programs relies on the substitution mechanism employed in the resolution. This aspect deserves a preliminary discussion. Traditionally, logic programming has both the theorem-proving and problem-solving perspectives [KP18b]. From the problem-solving perspective, the aim is to find a refutation of the goal  $\leftarrow G$ , which amounts to finding a proof of *some substitution instance* of  $G$ . From the theorem-proving perspective, the aim is to search for a proof of the goal  $G$  itself as an atom. The main difference is in the substitution mechanism of resolution: unification for the problem-solving and term-matching for the theorem-proving perspective. We will first explore computability within the theorem-proving perspective. As resolution therein is by term-matching, the probability  $\Pr_{\mathbb{P}}^{\text{TM}}(A)$  of proving a goal  $A$  in a PLP program  $\mathbb{P}$  is formulated as  $\Pr_{\mathbb{P}}^{\text{TM}}(A) := \sum_{|\mathbb{P}| \supseteq \mathbb{L} \Rightarrow A} \mu_{\mathbb{P}}(\mathbb{L})$ , where  $\mathbb{L} \Rightarrow A$  means that  $A$  is derivable in the sub-program

$\mathbb{L}$  (not to be confused with  $\mathbb{L} \vdash A$ , which stands for *some substitution instance* of  $A$  being derivable in  $\mathbb{L}$ , see (2.1)).

In our coalgebraic framework, the distribution semantics for general PLP programs is represented on ‘saturated’ trees, in which computations are performed by unification. However, following [BZ15], one can define the *TM (Term Matching) distribution tree* of a goal  $A$  in a program  $\mathbb{P}$  by ‘desaturation’ of the saturated distribution tree for  $A$  in  $\mathbb{P}$ . The coalgebraic definition, for which we refer to [BZ15], applies pointwise on the saturated tree the counit  $\epsilon_{\mathcal{U}\text{At}}: \mathcal{U}\mathcal{K}\mathcal{U}\text{At} \rightarrow \mathcal{U}\text{At}$  of the adjunction  $\mathcal{U} \dashv \mathcal{K}$  (cf. (4.3)). The TM distribution tree which results from ‘desaturation’ can be described very simply: at each layer of the starting saturated distribution tree, one prunes all the subtrees which are not labelled with the identity substitution  $id := x_1 \mapsto x_1, x_2 \mapsto x_2, \dots$ . In this way, the only remaining computation are those in which resolution only applies a non-trivial substitution on the clause side, that is, in which unification is restricted to term-matching.

**Computability of term-matching distribution semantics.** One may compute the success probability  $\Pr_{\mathbb{P}}^{\text{TM}}(A)$  in  $\mathbb{P}$  from the TM distribution tree of  $A$  in  $\mathbb{P}$ . The computation goes similarly to Algorithm 2: the problem amounts to calculating the probabilities of those deterministic subtrees of the distribution tree which prove the goal. We confine ourselves to some remarks on the aspects that require extra care, compared to the ground case.

- (1) The probability  $\Pr_{\mathbb{P}}^{\text{TM}}(A)$  is not computable in whole generality. It depends on whether one can decide all the proofs of  $A$  in the pure logic program  $|\mathbb{P}|$ , and there are various heuristics in logic programming for this task.
- (2) It is still possible to decide whether a subtree is deterministic, but the algorithm in the general case is a bit subtler, as it is now possible that two different goals match the same clause (instantiated in two different ways).
- (3) When calculating the probability of a deterministic subtree in the TM distribution tree, multiple appearances of a single clause (possibly instantiated with different substitutions) should be counted only once. In order to ensure this one needs to be able to identify which clause is applied at each step of the computation described by the distribution trees: this is precisely the reason of the addition of the clause labels in the coalgebra type of these trees, as discussed in Section 4.3.

We conclude by briefly discussing the problem-solving perspective, in which resolution is based on arbitrary unification rather than just term-matching. In standard SLD-resolution, computability relies on the possibility of identifying the *most general* unifier between a goal

and the head of a given clause. This can be done also within saturated distribution trees, since saturation supplies *all* the unifiers, thus in particular the most general one. This means that, in principle, one may compute the distribution semantics based on most general unification from the saturated distributed tree associated with a goal, with similar caveats as the ones we described for the term-matching case. However, the lack of a satisfactory coalgebraic treatment of most general unifiers [BZ15] makes us prefer the theorem-proving perspective discussed above, for which desaturation provides an elegant categorical formalisation.

### APPENDIX C. COMPUTABILITY OF THE WEIGHT SEMANTICS

In this appendix we will show how the weight  $\omega(A)$  of a goal  $A$  is computable from its derivation semantics, i.e. the weighted derivation tree  $\llbracket A \rrbracket_w$ . We focus on the ground case, and then conclude with a discussion on how the generalisation to the variable case goes.

One preliminary consideration concerns cycles in WLP. Indeed, in a weighted derivation tree it may be the case that an atom-node has a directed path to another atom-node labelled with the same atom. While different ways to assign meaning to cycles may be justifiable (for instance in the context of (co)inductive logic programming), the standard choice (see e.g. [CSS10]) is to regard any proof subtree containing a cycle as failing to prove the given goal. In concrete, this means that proof subtrees with cycles may be simply discarded when calculating the weight of the goal.

To make this precise in our computation, recall that a proof subtree for a goal  $A$  is a subtree of  $\llbracket A \rrbracket_w$  where we select exactly one child for each atom-node with children. We may define the weight of a proof subtree for  $A$  as the weight of  $A$  in that particular tree, calculated using (5.1). Note the overall weight of a goal  $A$  in  $\llbracket A \rrbracket_w$  is then the  $\oplus$ -sum of the weights of its proof subtrees.

This means that, from a computational viewpoint, discarding a proof subtree is equivalent to assigning weight  $\mathbf{0}$  (the  $\oplus$ -unit) to that proof subtree. In order to assign weight  $\mathbf{0}$  to each proof subtree with cycles, we look for the first appearance of a cycle, re-label the next edge with  $\mathbf{0}$ , and discard all the descendants. This works because, according to (5.1) and the fact that  $\mathbf{0} \otimes c = c \otimes \mathbf{0} = \mathbf{0}$  for any  $c \in K$ , any finite proof subtree containing an edge labelled with  $\mathbf{0}$  has weight  $\mathbf{0}$ . We illustrate this idea with the following example.

**Example C.1.** Recall  $\mathbb{P}^{gsp}$  in Example 5.4, and consider the goal  $\text{reachable}(c)$ . Its weighted derivation tree  $\llbracket \text{reachable}(c) \rrbracket_{\mathbb{P}^{gsp}}$  includes a finite path (subtree) witnessing a proof of the goal.

$$\text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{0} \bullet \rightarrow \text{init}(a) \xrightarrow{0} \bullet$$

Intuitively, this says that  $c$  is reachable from the initial state  $a$  in (5.2), with weight 4. However, in  $\llbracket \text{reachable}(c) \rrbracket_{\mathbb{P}^{gsp}}$  there are also other proof subtrees, which feature cycles, as for instance

$$\text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{9} \bullet \rightarrow \text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{0} \bullet \rightarrow \text{init}(a) \xrightarrow{0} \bullet \quad (\text{C.1})$$

$$\text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{9} \bullet \rightarrow \text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{9} \bullet \rightarrow \text{reach}(c) \rightarrow \dots \quad (\text{C.2})$$

Our algorithm for computing the weight of  $w(\text{reachable}(c))$  will prune both (C.1) and (C.2), letting them both become equal to

$$\text{reach}(c) \xrightarrow{4} \bullet \rightarrow \text{reach}(a) \xrightarrow{+\infty} \bullet \quad (\text{C.3})$$

where  $+\infty$  is the element  $\mathbf{0}$  in the semiring of  $\mathbb{P}^{gsp}$ . To see that this procedure yields the correct result, note that the weight of the proof subtree (C.3) is  $4 + (+\infty) = +\infty$ , which does not affect the calculation of the weight ( $= 4$ ) of  $\text{reachable}(c)$ .

Now we give the generic algorithms which compute the weight of a goal  $A$  given its weighted derivation tree  $\llbracket A \rrbracket_w$  in the program  $\mathbb{W}$ . For readability, we divide our approach into 2 algorithms.

---

**Algorithm 3** Prune a weighted derivation tree
 

---

**Input:** a weighted derivation tree  $\mathcal{T}$

**Output:**  $\mathcal{T}$  pruned all cycles

```

1: for path  $\pi$  in  $\mathcal{T}$  do
2:   atom_list = [ ]
3:   for node  $v$  in  $\pi$  do
4:     if label( $v$ ) in atom_list then
5:        $t =$  grandparent of  $v$ ,  $u =$  parent of  $v$ 
6:       cut the descendants of  $u$  from  $\mathcal{T}$ 
7:       re-label the edge  $t \rightarrow u$  by  $\mathbf{0}$ 
8:     else
9:       atom_list.append(label( $v$ ))
10: return tree  $\mathcal{T}$ 

```

---



---

**Algorithm 4** Calculate weight of the root node from a pruned weighted derivation tree
 

---

**Input:** the weighted derivation tree  $\mathcal{T}$  for the goal  $A$

**Output:** the weight of  $A$

```

1: Prune  $\mathcal{T}$  using Algorithm 3
2: function WEIGHT( $x, \mathcal{T}$ )
3:   sum_list = [ ]
4:   for children  $y$  of  $x$  do
5:     prod_list = [label( $x \rightarrow y$ )]            $\triangleright$  first get the weight of the clause
6:     for children  $z$  of  $y$  do
7:       prod_list.append(WEIGHT( $z, \mathcal{T}$ ))
8:       cur_weight =  $\otimes$  prod_list            $\triangleright$  calculate the  $\otimes$ -product of all the values in
prod_list
9:     sum_list.append(cur_weight)
10:  return  $\oplus$  sum_list            $\triangleright$  return the  $\oplus$ -sum of all the values in product_list
11:  $s =$  root of  $\mathcal{T}$             $\triangleright$  calculate the weight of node  $s$  in  $\mathcal{T}$ 
12: return WEIGHT( $s, \mathcal{T}$ )

```

---

Algorithm 3 uses depth-first search to find the first node whose atom label already appeared in some of its ancestors. This search procedure terminates because the set  $\mathbf{At}$  of atoms is finite. Since a weighted derivation tree is finitely branching, the whole algorithm also terminates.



Algorithm 4 is a typical divide-and-conquer algorithm, where the weight of an atom-node is calculated via the weights of its (atom-)grandchildren. The procedure is simply spelling out (5.1) in the weighted derivation tree.

We conclude by sketching how this algorithm generalises to one for general WLP programs. First, recall the theorem-proving and problem-solving perspectives on logic programming, as discussed in Appendix B. For WLP, the theorem-proving aspect is prevalent in most applications (for example [CSS10, EB20, DEKM98]). Within this perspective, the task is to compute the weight of the goal itself, rather than the weight of some substitution instance of the goal. To compute such term-matching weight semantics, say for a goal  $A$ , one should first apply desaturation (as in Appendix B, see also [BZ15]) to the weighted saturated derivation tree  $\llbracket A \rrbracket_{\text{u}\sharp}$ , obtaining a new tree  $T_A$ . Concretely,  $T_A$  is obtained from  $\llbracket A \rrbracket_{\text{u}\sharp}$  by pruning any subtree which is rooted by substitutions other than the identity substitution.

Second, one may tweak Algorithm 3 and 4 for ground WLP in order to deal with the substitution-nodes in ‘desaturated’ weighted derivation trees. The term-matching weight semantics of  $A$  can then be computed by applying the modified algorithms to  $T_A$ . Note that to guarantee termination of the computation, we require that for every atom  $A$ , there are only finitely many grounding  $H\tau \leftarrow B_1\tau, \dots, B_k\tau$  of the clauses whose heads  $H\tau$  is  $A$ .

#### APPENDIX D. MISSING PROOFS

**Proposition D.1.** *The operation  $\text{pw} : \mathcal{M}_{pr} \Rightarrow \mathcal{D}_{\leq 1}\mathcal{P}_f$  defined in Definition 3.12 is a natural transformation.*

*Proof.* The naturality for  $\text{pw}$  boils down to showing that for arbitrary  $h : X \rightarrow Y$ ,  $\mathcal{D}_{\leq 1}\mathcal{P}_f(h) \circ \text{pw}_X = \text{pw}_Y \circ \mathcal{M}_{pr}(h) : \mathcal{M}_{pr}(X) \rightarrow \mathcal{D}_{\leq 1}\mathcal{P}_f(Y)$ . In this proof we fix an arbitrary function  $h : X \rightarrow Y$ . We start from an arbitrary  $\varphi \in \mathcal{M}_{pr}(X)$ . The case where  $\text{supp}(\varphi) = \emptyset$  is easy. So we assume that  $\varphi$  has a finite non-empty support  $\text{supp}(\varphi)$ . On one hand,  $\text{pw}_X(\varphi)$  has support  $\mathcal{P}(\text{supp}(\varphi))$ , and for any  $A \in \mathcal{P}(\text{supp}(\varphi))$ ,  $\text{pw}_X(\varphi)(A) = \left( \prod_{a \in A} \varphi(a) \right) \cdot \left( \prod_{a \in \text{supp}(\varphi) \setminus A} (1 - \varphi(a)) \right)$ . Then the support of  $\mathcal{D}_{\leq 1}\mathcal{P}_f(h)(\text{pw}_X(\varphi))$  is  $h[\text{supp}(\text{pw}_X(\varphi))] = h[\mathcal{P}(\text{supp}(\varphi))]$ , and for each  $B \in h[\mathcal{P}(\text{supp}(\varphi))]$ ,

$$\begin{aligned} \mathcal{D}_{\leq 1}\mathcal{P}_f(h)(\text{pw}_X(\varphi))(B) &= \sum_{\substack{A \in \text{supp}(\text{pw}_X(\varphi)) \\ h[A]=B}} \text{pw}_X(\varphi)(A) \\ &= \sum_{\substack{A \in \mathcal{P}(\text{supp}(\varphi)) \\ h[A]=B}} \text{pw}_X(\varphi)(A) \\ &= \sum_{\substack{A \in \mathcal{P}(\text{supp}(\varphi)) \\ h[A]=B}} \left( \prod_{a \in A} \varphi(a) \cdot \prod_{a \in \text{supp}(\varphi) \setminus A} (1 - \varphi(a)) \right) \end{aligned} \quad (\text{D.1})$$

On the other hand,  $\mathcal{M}_{pr}(h)(\varphi)$  has support  $h[\text{supp}(\varphi)]$ , and for arbitrary  $y \in h[\text{supp}(\varphi)]$ ,  $\mathcal{M}_{pr}(h)(\varphi)(y) = \bigvee_{x \in h^{-1}(y)} \varphi(x) = 1 - \prod_{x \in h^{-1}(y)} (1 - \varphi(x))$ . Then  $\text{pw}_Y(\mathcal{M}_{pr}(h)(\varphi))$  has support  $\mathcal{P}(\text{supp}(\mathcal{M}_{pr}(h)(\varphi))) = \mathcal{P}(h[\text{supp}(\varphi)])$ . For each  $B' \in \mathcal{P}(h[\text{supp}(\varphi)])$ ,

$$\text{pw}_Y(\mathcal{M}_{pr}(h)(\varphi))(B') = \prod_{y \in B'} \mathcal{M}_{pr}(h)(\varphi)(y) \cdot \prod_{y \in h[\text{supp}(\varphi)] \setminus B'} (1 - \mathcal{M}_{pr}(h)(\varphi)(y)) \quad (\text{D.2})$$

Note that the two supports  $h[\mathcal{P}(\text{supp}(\varphi))]$  and  $\mathcal{P}(h[\text{supp}(\varphi)])$  are equivalent:  $B \in h[\mathcal{P}(\text{dom}(\varphi))]$  if and only if  $\exists A \subseteq \text{dom}(\varphi)$  such that  $h[A] = B$ ;  $B' \in \mathcal{P}(h[\text{supp}(\varphi)])$  if and only if  $B' \subseteq h[\text{supp}(\varphi)]$  if and only if  $\exists A' \subseteq \text{supp}(\varphi)$  such that  $h[A'] = B'$ .

We spell out the operation of  $\mathcal{M}_{pr}$  on  $h$  in (D.2), and the steps (D.3), (D.4) below will be explained in detail after the calculation:

$$\begin{aligned} & \text{pw}_Y(\mathcal{M}_{pr}(h)(\varphi))(B') \\ &= \prod_{y \in B'} \left( 1 - \prod_{x \in h^{-1}(y)} (1 - \varphi(x)) \right) \cdot \prod_{y \in h[\text{supp}(\varphi)] \setminus B'} \left( 1 - \left( 1 - \prod_{x \in h^{-1}(y)} (1 - \varphi(x)) \right) \right) \\ &= \prod_{y \in B'} \left( 1 - \prod_{x \in h^{-1}(y)} (1 - \varphi(x)) \right) \cdot \prod_{y \in h[\text{supp}(\varphi)] \setminus B'} \prod_{x \in h^{-1}(y)} (1 - \varphi(x)) \\ &\stackrel{(\text{D.5})}{=} \prod_{y \in B'} \left( \sum_{Z \subseteq_i h^{-1}(y)} \left( \prod_{x \in Z} \varphi(x) \cdot \prod_{x \in h^{-1}(y) \setminus Z} (1 - \varphi(x)) \right) \right) \cdot \prod_{\substack{y \in h[\text{supp}(\varphi)] \setminus B' \\ x \in h^{-1}(y)}} (1 - \varphi(x)) \quad (\text{D.3}) \end{aligned}$$

$$= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{x \in A'} \varphi(x) \cdot \prod_{x \in h^{-1}[B'] \setminus A'} (1 - \varphi(x)) \right) \cdot \prod_{\substack{y \in h[\text{supp}(\varphi)] \setminus B' \\ x \in h^{-1}(y)}} (1 - \varphi(x)) \quad (\text{D.4})$$

$$= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{x \in A'} \varphi(x) \cdot \prod_{x \in h^{-1}[B'] \setminus A'} (1 - \varphi(x)) \right) \cdot \prod_{\substack{x \in \text{supp}(\varphi) \\ x \notin h^{-1}[B']}} (1 - \varphi(x))$$

$$= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{x \in A'} \varphi(x) \cdot \prod_{x \in h^{-1}[B'] \setminus A'} (1 - \varphi(x)) \cdot \prod_{\substack{x \in \text{supp}(\varphi) \\ x \notin h^{-1}[B']}} (1 - \varphi(x)) \right)$$

$$= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{x \in A'} \varphi(x) \cdot \prod_{x \in \text{supp}(\varphi) \setminus A'} (1 - \varphi(x)) \right)$$

$$\stackrel{(\text{D.1})}{=} \mathcal{D}_{\leq 1} \mathcal{P}_f(h)(\mathcal{M}_{pr}(\varphi))(B')$$

In (D.3),  $\subseteq_i$  is the relation  $\subseteq$  restricted to non-empty subsets. Formally, (D.3) is obtained via the following equation, for arbitrary set  $X$ ,  $\psi \in \mathcal{M}_{pr}(X)$  and set  $A \subseteq \text{supp}(\psi)$ :

$$1 - \prod_{a \in A} (1 - \psi(a)) = \sum_{Z \subseteq_i A} \left( \prod_{a \in Z} \psi(a) \cdot \prod_{a \in A \setminus Z} (1 - \psi(a)) \right) \quad (\text{D.5})$$

Intuitively, one can think of  $A$  as a set of independent events, and  $\psi(a)$  as the probability of the event  $a \in A$ . Then the left-hand-side of (D.5) calculates 1 minus the probability that none of  $a \in A$  happens; the right-hand-side sums up the probabilities of all possible worlds in which some of  $a \in A$  happens. Both calculate the same value, namely the probability that at least one  $a \in A$  happens.

The expression (D.4) is obtained by expanding  $\prod_{y \in B'} \sum_{Z \subseteq_i h^{-1}(y)} G(y, Z)$  into a summation, where  $G(y, Z) = \prod_{x \in Z} \varphi(x) \cdot \prod_{x \in h^{-1}(y) \setminus Z} (1 - \varphi(x))$ . In the resulting summation, each summand chooses exactly one element from each  $\sum_{Z \subseteq_i h^{-1}(y)} G(y, Z)$ , for all  $y \in B'$ . Moreover, every such choice corresponds to exactly one  $A'$  satisfying  $h[A'] = B'$ . So we have

$$\begin{aligned}
\prod_{y \in B'} \sum_{Z \subseteq_i h^{-1}(y)} G(y, Z) &= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \prod_{y \in B'} G(y, h^{-1}(y) \cap A') \\
&= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{y \in B'} \left( \prod_{x \in h^{-1}(y) \cap A'} \varphi(x) \cdot \prod_{x \in h^{-1}(y) \setminus (h^{-1}(y) \cap A')} (1 - \varphi(x)) \right) \right) \\
&= \sum_{\substack{A' \subseteq \text{supp}(\varphi) \\ h[A'] = B'}} \left( \prod_{x \in A'} \varphi(x) \cdot \prod_{x \in h^{-1}[B'] \setminus A'} (1 - \varphi(x)) \right)
\end{aligned}$$

and this is exactly the summation in (D.4).

Therefore  $(\mathcal{D}_{\leq 1} \mathcal{P}_f(h) \circ \text{pw}_X)(\varphi)$  and  $(\text{pw}_Y \circ \mathcal{M}_{pr}(h))(\varphi)$  have the same support, and both have the same values for arbitrary  $B$  in the support.  $\square$