

# SMProbLog: Stable Model Semantics in ProbLog for Probabilistic Argumentation<sup>†</sup>

PIETRO TOTIS, LUC DE RAEDT<sup>1</sup> and ANGELIKA KIMMIG<sup>1</sup>  
*KU Leuven, Department of Computer Science, Leuven.AI, B-3000 Leuven, Belgium*  
(e-mails: [pietro.totis@cs.kuleuven.be](mailto:pietro.totis@cs.kuleuven.be), [luc.deraedt@cs.kuleuven.be](mailto:luc.deraedt@cs.kuleuven.be),  
[angelika.kimmig@cs.kuleuven.be](mailto:angelika.kimmig@cs.kuleuven.be))

*submitted 23 November 2021; revised 2 March 2023; accepted 6 April 2023*

---

## Abstract

Argumentation problems are concerned with determining the acceptability of a set of arguments from their relational structure. When the available information is uncertain, probabilistic argumentation frameworks provide modeling tools to account for it. The first contribution of this paper is a novel interpretation of probabilistic argumentation frameworks as probabilistic logic programs. Probabilistic logic programs are logic programs in which some of the facts are annotated with probabilities. We show that the programs representing probabilistic argumentation frameworks do not satisfy a common assumption in probabilistic logic programming (PLP) semantics, which is, that probabilistic facts fully capture the uncertainty in the domain under investigation. The second contribution of this paper is then a novel PLP semantics for programs where a choice of probabilistic facts does not uniquely determine the truth assignment of the logical atoms. The third contribution of this paper is the implementation of a PLP system supporting this semantics: SMProbLog. SMProbLog is a novel PLP framework based on the PLP language ProbLog. SMProbLog supports many inference and learning tasks typical of PLP, which, together with our first contribution, provide novel reasoning tools for probabilistic argumentation. We evaluate our approach with experiments analyzing the computational cost of the proposed algorithms and their application to a dataset of argumentation problems.

**KEYWORDS:** probabilistic logic programming, ProbLog, distribution semantics, stable model semantics, probabilistic argumentation

---

## 1 Introduction

Designing systems that are able to argue and persuade is a very relevant and challenging task in artificial intelligence. In real-world scenarios the information available is often incomplete or questionable, therefore the ability to take into account uncertainty is fundamental in such situations. Probabilistic logic programming (PLP) frameworks and languages, such as PRISM (Sato and Kameya 2008), ICL (Poole 2008), ProbLog (De Raedt *et al.* 2007), or LPAD/CP-logic (Vennekens *et al.* 2009), are designed to provide powerful

<sup>1</sup> Shared last authors.

<sup>†</sup> This work was supported by the FWO project N. G066818N and the Flanders AI program.



Table 1. *Argumentation frameworks overview*

Framework	Epistemic	PLP	EM learning
MetaProblog (Mantadelis and Bistarelli 2020)	✗	✓	✗
Epistemic graphs (Hunter <i>et al.</i> 2020)	✓	✗	✗
SMProbLog	✓	✓	✓

general-purpose tools for modeling and reasoning about structured, uncertain domains. On the other hand, abstract argumentation frameworks (Dung 1995) aim specifically at describing argumentation processes and reasoning about the set of acceptable arguments. Consider, for instance, the following example of argumentative microtext adapted from (Stede *et al.* 2016):

*Example 1*

Yes, it is annoying and cumbersome to separate your rubbish properly all the time ( $a_1$ ), but small gestures become natural by daily repetition ( $a_2$ ). Three different bin bags stink away in the kitchen and have to be sorted into different wheelie bins ( $a_3$ ). But still Germany produces way too much rubbish ( $a_4$ ) and too many resources are lost when what actually should be separated and recycled is burnt ( $a_5$ ). We Berliners should take the chance and become pioneers in waste separation! ( $a_6$ )

An argumentation problem defines arguments, for example,  $a_1, a_2, \dots$  and their relations, for example  $a_1$  attacks  $a_2$ ,  $a_3$  supports  $a_1 \dots$  from which a set of acceptable arguments is derived according to a given semantics.

When uncertainty is part of the argumentation problem, we want to reason about it in order to answer queries like “What is the likelihood of accepting argument  $a_6$ ?”. Inferring the probability of one variable taking a particular value is a typical PLP task. PLP frameworks offer a wide range of tools and algorithms for probabilistic inference and learning on probabilistic logic programs. Probabilistic argumentation systems, on the other hand, propose different combinations of argumentation frameworks, probability interpretations and reasoning systems, tailored to manipulating probabilities in argumentation. When considering PLP and argumentation systems, the question is natural: can PLP effectively model argumentation processes and reason over its intrinsic uncertainty?

This question has been only partially investigated so far: of the two main interpretations of probabilities in argument graphs (Hunter 2013), namely the *constellations* approach and the *epistemic* approach, only the former has been studied in the context of PLP (Mantadelis and Bistarelli 2020). This paper fills this gap and shows how PLP can be used to reason about a probabilistic argumentation graph with an epistemic view of its probabilities. Table 1 summarizes the differences in the approaches.

In the constellations approach, probabilities describe a probability distribution over argument (sub)graphs. Such a distribution is used to define the probability of an argument being in the set of accepted arguments according to the argumentation semantics of choice. On the contrary, in the epistemic approach probabilities are a direct measure of the belief of an agent in an argument, and therefore of its acceptability. Previous work on probabilistic argument graphs under the epistemic approach focuses on model-

Table 2. *PLP frameworks overview*

Framework	Model semantics	Distribution semantics
PLP frameworks (ProbLog, PRISM, ...)	Least or 2-valued Well-founded model	✓
Probabilistic ASP (P-Log, LP <sup>MLN</sup> )	Stable model	✗
SMProbLog	Stable model	✓

ing a family of probability distributions that are compatible with the graph and given constraints (Hunter et al. 2020).

Instead of considering a probabilistic argument graph as a model for a family of probability distributions over fixed beliefs that have to satisfy constraints, we regard a probabilistic argument graph as a model for a single joint probability distribution over arguments. We view nodes as random variables associated to the arguments, whose prior is the epistemic arbitrary assignment of probabilities to arguments. We then interpret their relations as the conditional (in)dependencies influencing the posterior joint probability (belief) distribution of the random variables. Given a set of independent beliefs regarding arguments and relations, modeled by means of PLP, we thus define the joint probability (belief) of the arguments by means of probabilistic inference.

The joint distribution defines *marginal and conditional* beliefs that are coherent with the given priors and argumentative (logic) structure. The advantage of this method is that the joint belief distribution allows us to answer a broad range of questions that are of interest in the argumentation setting. For example, the marginal probability of (belief in) an argument (“What is the belief in  $a_1$  considering the influence of the other arguments?”), relating the beliefs of two arguments (“How does my belief in  $a_1$  change if  $a_5$  is accepted?”), or perform typical PLP tasks such as learning the beliefs of an agent given a set of observations of accepted arguments.

We propose a new PLP framework where this is possible, since we argue that existing PLP systems are limited in what problems they can model and reason about. In fact, traditional PLP frameworks do not allow modeling cyclic relations involving negations. This is a pattern often found in argument graphs, where reciprocal attacks are common. For example, we will encode this pattern, that is, accepting  $a_1$  inhibits my belief in  $a_2$  and vice versa, with the logic rules  $\neg a_1 \leftarrow a_2$ .  $\neg a_2 \leftarrow a_1$ . A model containing such rules would not be valid in traditional PLP frameworks, because they assume that each (deterministic) logic program induced by the probabilistic model has exactly one two-valued well-founded model (Gelder et al. 1991). Programs containing cyclic dependencies involving negation usually do not satisfy this assumption. Probabilistic answer set programming (ASP) frameworks, for example, Baral et al. (2009) or Lee and Wang (2016), do not make this assumption, however the definition of the probability distribution is not based on distribution semantics (Sato 1995), hence there are differences that we will analyze throughout the paper. Table 2 summarizes the main characteristics of PLP frameworks.

In Section 3, we introduce a framework modeling probabilistic argumentation problems by means of typical PLP modeling techniques. In Section 5, we analyze the properties

of this approach. In Sections 4 and 6, we propose a PLP system, SMProbLog, based on a new semantics, where it is possible to reason over such models. The advantages of reasoning over probabilistic argument graphs by means of PLP are:

- *Succinctness and expressivity*: expressing complex interactions between random variables with the simplicity of (first-order) logic programming rules.
- *Flexibility and modularity*: improving a model for a specific application domain without having to modify neither the language nor the reasoning algorithms.
- *PLP tools*: applying to argumentation problems a general suite of PLP inference and learning algorithms to any model within a broad class.

To summarize, the key contributions of the paper are:

1. We define a novel semantics for PLP based on stable model semantics (Gelfond and Lifschitz 1988) (Section 4).
2. We develop an implementation of a PLP framework, SMProbLog<sup>2</sup>, derived from ProbLog2 (Dries *et al.* 2015) that supports inference and learning tasks under the new semantics (Section 6).
3. We show an application of SMProbLog to encode and solve probabilistic argumentation problems in a novel reasoning framework based on an epistemic interpretation of probability (Sections 3 and 5).

## 2 Background

In this paper we present a PLP framework, SMProbLog based on ProbLog2: we describe the relevant background about syntax, the semantics for the probabilistic component of ProbLog programs, and the semantics for the models of a logic program we will consider throughout the paper. We also describe the corresponding inference and learning tasks under the new semantics. Finally, we also consider the application of SMProbLog to probabilistic argumentation problems, hence we discuss the required background in argumentation.

### 2.1 Probabilistic logic programming

*ProbLog.* ProbLog (De Raedt *et al.* 2007) is a probabilistic language extending Prolog, where facts and clauses are annotated with (mutually independent) probabilities. Probabilistic logic programs based on distribution semantics, such as ProbLog, can be viewed as a “programming language” generalization of Bayesian networks (Sato 1995). A *normal* logic program  $\mathcal{L}$  is a finite set of normal rules of the form  $h \leftarrow b_1, \dots, b_n$ . We use the standard terminology about atoms, terms, predicates, and literals. The *head*  $h$  is an atom and the *body*  $b_1, \dots, b_n$  is a logical conjunction of literals. Rules with empty body are called facts. We denote with  $\sim a$  the negation as failure of an atom  $a$ , used to make literals, and with  $\neg a$  the classical logical negation of an atom  $a$ .

We model uncertainty in logic programs by annotating facts with probabilities: the probabilities of facts are mutually independent and each *probabilistic fact*, that is,  $p :: f.$ ,

<sup>2</sup> <https://github.com/PietroTotis/smProblog>

corresponds to an atomic choice, that is, a choice between including  $f$  in the program (with probability  $p$ ) or discarding it ( $1 - p$ ). A *probabilistic normal logic program*  $\mathcal{L} = F \cup R$  is thus a set  $F$  of facts annotated with probabilities plus a set of logic rules  $R$ . The set of facts  $F$  is disjoint from the heads of the rules  $R$ , nonetheless rule heads can be annotated as syntactic sugar:  $p :: h \leftarrow b_1, \dots, b_n$  is equivalent to a new fact  $p :: f$  plus  $h \leftarrow f, b_1, \dots, b_n$ .

### Example 2

The well-known alarm Bayesian network (Pearl 1989) can be encoded as a ProbLog program. An alarm is triggered by either an earthquake or a burglary: if the neighbor is at home you will receive a call. The neighbor at home, the earthquake, and the burglary are independent events occurring with a certain probability, thus encoded as probabilistic facts. Probabilistic facts are marginally independent and therefore correspond to nodes of the Bayesian network that do not have common parents or ancestors. The logic rules determine whether you get a call or not, and thus correspond to the conditional dependencies of the Bayesian network.

---

```

0.1 :: burglary.
0.2 :: earthquake.
0.5 :: neighbor_at_home.
alarm ← earthquake.
alarm ← burglary.
neighbor_calls ← alarm, neighbor_at_home.
query(alarm).
query(neighbor_calls).

```

---

*Distribution semantics.* A probabilistic logic program  $\mathcal{L} = F \cup R$  is queried for the likelihood of atoms. The probability of an atom is commonly defined by the *distribution semantics* (Sato 1995). A *total choice*  $\omega$  is a combination of atomic choices over all probabilistic facts, that is, a subset of the set of all ground facts  $F$ , that is,  $\omega \subseteq F$ . We denote the set of all total choices of  $\mathcal{L}$ , that is, the power set of  $F$ , with  $\Omega_{\mathcal{L}}$ . The probabilities of the facts define a probability distribution over the  $2^{|F|}$  total choices. The possible non-probabilistic subprograms  $\omega \cup R, \omega \in \Omega_{\mathcal{L}}$ , which we call *possible worlds*, are obtained from  $\mathcal{L}$  by including or discarding a fact according to the corresponding atomic choice. An interpretation is an assignment of a truth value to all atoms in the Herbrand base of a program. An interpretation is called a model of the theory if it satisfies all formulas in the theory.

### Example 3

Consider the alarm example (Example 2), there are  $2^3=8$  possible worlds (listed in Example 4), corresponding to the combinations of choices for the probabilistic facts *earthquake*, *burglary*, *neighbor\_at\_home*. Consider the possible world where we discard *earthquake* and include the other two: then the alarm will be triggered, that is, *alarm* is inferred from *burglary*, and thus the neighbor calls, because both *alarm* and *neighbor\_at\_home* are true. The model for the possible world denoted by

$\{\text{burglary}, \text{neighbor\_at\_home}\}$  is thus  $\{\text{burglary}, \text{neighbor\_at\_home}, \text{alarm}, \text{neighbor\_calls}\}$ .

The distribution semantics relies on a one-to-one mapping between possible worlds and models: many PLP frameworks, for example, ProbLog, PRISM, and CP-Logic, restrict the class of valid inputs to programs where each total choice corresponds to a two-valued well-founded model (Gelder *et al.* 1991). The underlying assumption is that choices are modeled exclusively by means of probabilistic facts, and that they are independent. Under this assumption many algorithms and reasoning techniques have been developed to perform inference and parameter learning tasks, which we present in the rest of this section.

*Inference.* The MARG inference task is the task of computing the probability of success of a query, that is, the sum of the probability of each possible world where the query is true, under the (possibly empty) given evidence. Evidence is a set of atoms whose truth value is known. In particular:

*Definition 1 (The MARG Inference task)*

**Given**

- A program  $\mathcal{L}$ : let  $G$  be the set of all ground (probabilistic and derived) atoms of  $\mathcal{L}$ .
- A set  $E \subseteq G$  of observed atoms (evidence), along with a vector  $e$  of corresponding observed truth values ( $E = e$ ).
- A set  $Q \subseteq G$  of atoms of interest (queries).

**Find** the marginal distribution of every query atom given the evidence, that is, computing  $P(q | E = e)$  for each  $q \in Q$ .

In the ProbLog2 system the probabilistic inference task is reduced to a *weighted model counting problem (WMC)* (Cadoli and Donini 1997). WMC is the problem of computing the weight of a propositional logic formula  $\varphi$ , given a weight function  $w$  that assigns a positive (real) weight value to each literal  $v$  and  $\neg v$  ( $v \in G$ ). Let  $MOD(\mathcal{L})$  be the set of valid interpretations for  $\mathcal{L}$ , the *WMC* of  $\varphi$  is:

$$WMC_{\mathcal{L}}(\varphi) = \sum_{M \in MOD(\mathcal{L}), M \models \varphi} \prod_{l \in M} w(l).$$

The probability  $P(q | E = e)$  is computed as  $WMC_{\mathcal{L}}(\varphi)$  where  $\varphi$  is the propositional representation of  $q \wedge E = e$ . The weight function is defined as follows: for all probabilistic facts ( $p : f$ ),  $w(f) = p$  and  $w(\neg f) = 1 - p$ ; for all atoms  $a \in E$  if  $a$  is true (false)  $w(a) = 1$ ,  $w(\neg a) = 0$  (resp.  $w(a) = 0$ ,  $w(\neg a) = 1$ ); for all remaining (logical) atoms  $a$ ,  $w(a) = w(\neg a) = 1$ .

*Example 4*

The possible world of Example 2 corresponding to  $\{\text{burglary}, \text{neighbor\_at\_home}\}$  has probability  $0.1 \cdot (1 - 0.2) \cdot 0.5 = 0.04$ . To find the probability of being called we sum the probabilities of all other possible worlds whose model contains *neighbor\_calls*, namely  $\{\text{earthquake}, \text{neighbor\_at\_home}\}$  and  $\{\text{burglary}, \text{earthquake}, \text{neighbor\_at\_home}\}$ , which are respectively  $(1 - 0.1) \cdot 0.2 \cdot 0.5 = 0.09$  and  $0.1 \cdot 0.2 \cdot 0.5 = 0.01$ . Therefore,  $P(\text{neighbor\_calls}) = WMC(\text{neighbor\_calls}) = 0.04 + 0.09 + 0.01 = 0.14$ .

Possible world's choices $\omega$	Model	$P(\omega)$
$\{\}$	$\omega$	0.36
$\{neighbor\_at\_home\}$	$\omega$	0.36
$\{burglary\}$	$\omega \cup \{alarm\}$	0.04
$\{earthquake\}$	$\omega \cup \{alarm\}$	0.09
$\{neighbor\_at\_home, burglary\}$	$\omega \cup \{alarm, neighbor\_calls\}$	<b>0.04</b>
$\{burglary, earthquake\}$	$\omega \cup \{alarm\}$	0.01
$\{earthquake, neighbor\_at\_home\}$	$\omega \cup \{alarm, neighbor\_calls\}$	<b>0.09</b>
$\{burglary, earthquake, neighbor\_at\_home\}$	$\omega \cup \{alarm, neighbor\_calls\}$	<b>0.01</b>

The task of model counting is  $\#P$ -complete in general, therefore the logic program is transformed into a representation where this task becomes tractable (Fierens et al. 2011).  $\mathcal{L}$  is in fact transformed by means of knowledge compilation into a *smooth d-DNNF* formula (deterministic decomposable negation normal form) (Darwiche 2004) or a *sentential decision diagram* (Choi et al. 2013). In this paper we focus on d-DNNFs because we base our approach on a knowledge compiler to d-DNNFs for stable model counting (Aziz et al. 2015).

#### Definition 2

An *NNF* is a rooted directed acyclic graph in which each leaf node is labeled with a literal and each internal node is labeled with a disjunction or conjunction. A *smooth d-DNNF* is an NNF with the following properties:

- **Deterministic:** for all disjunctive nodes the children represent formulas pairwise inconsistent.
- **Decomposable:** the subtrees rooted in two children of a conjunction node do not have atoms in common.
- **Smooth:** all children of a disjunction node use the same set of atoms.

On d-DNNFs the task of model counting becomes tractable. The d-DNNF is further transformed into an equivalent arithmetic circuit, by replacing conjunctions and disjunctions respectively with multiplication and summation nodes, and by replacing leaves with the weight of the corresponding literals. Arithmetic circuits allow us to efficiently perform the task of *WMC*.

#### Example 5

Consider the rule  $neighbor\_calls \leftarrow alarm, neighbor\_at\_home$ . The corresponding smooth d-DNNF and arithmetic circuit are represented in Figure 1. The root of the circuit represents all possible assignments to the variables. The leftmost and-node corresponding to the assignment  $calls$  corresponds in the arithmetic circuit to the weight 0.14, the probability of it being true. Vice versa, the top and-node, corresponding to the assignment  $\neg calls$ , represents in the arithmetic circuit the weight 0.86. We simplified Example 2 by considering  $alarm$  a probabilistic fact with probability 0.28, which is the probability that can be inferred from  $burglary$  and  $earthquake$  (cfr. Example 4).

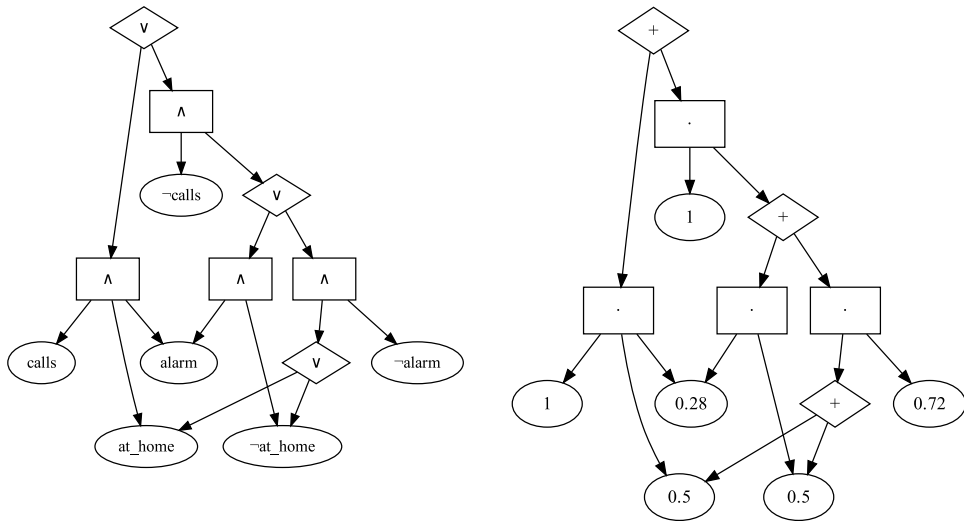


Fig. 1. Smooth d-DNNF for the implication  $neighbor\_calls \leftarrow alarm, neighbor\_at\_home$  (left) and the corresponding arithmetic circuit (right).

*Learning.* Learning from interpretations of parameters in a ProbLog program is implemented in a likelihood maximization setting (Gutmann et al. 2011). In the *learning from interpretations* task the input is a program  $\mathcal{L}$  where some probabilities are unknown (parameters), and a set of interpretations for  $\mathcal{L}$ . Interpretations can be partial, that is, the observation of truth values contains only a subset of all atoms. The goal is to estimate the value of the parameters such that the predicted values maximize the likelihood of the given interpretations:

*Definition 3 (Max-Likelihood Parameter Estimation)*

**Given**

- A program  $\mathcal{L}(\mathbf{p}) = F \cup R$  where  $F$  is a set of probabilistic facts and  $R$  contains the rules describing the background knowledge.  $\mathbf{p} = \langle p_1, \dots, p_N \rangle$  is a set of unknown parameters attached to probabilistic facts.
- A set  $\mathbb{I}$  of (partial) interpretations  $\{I_1, \dots, I_M\}$  as training examples.

**Find** the maximum likelihood probabilities  $\hat{\mathbf{p}} = \langle \hat{p}_1, \dots, \hat{p}_N \rangle$  for the interpretations in  $\mathbb{I}$ . Formally,

$$\hat{\mathbf{p}} = \arg \max_{\mathbf{p}} P(\mathbb{I} | \mathcal{L}(\mathbf{p})) = \arg \max_{\mathbf{p}} \prod_{m=1}^M P(I_m | \mathcal{L}(\mathbf{p})).$$

Parameters are iteratively updated by an alternation of an expectation step (E step) which computes the expected value of the learnable parameters under the current model, and the maximization step (M step) which updates the parameters estimates to maximize the likelihood of the observed data.

*Example 6*

Consider Example 5 but now the probability of *alarm* is a learnable parameter. In ProbLog the expression  $t(0.4) :: alarm$ . denotes a learnable probability which is initialized at 0.4 in the expectation-maximization (EM) algorithm:



---

$t(0.4) :: \text{alarm.}$   
 $0.5 :: \text{at\_home.}$   
 $\text{calls} \leftarrow \text{alarm, at\_home.}$

---

Given a set of interpretations reflecting the original probability distribution we can learn the probability of  $0.28 :: \text{alarm.}$  even if the observation of alarm is not given. In fact, we provide three types of partial interpretations:

$I_1: \{\neg \text{calls}, \neg \text{at\_home}\}$  (50 examples)

$I_2: \{\text{calls}, \text{at\_home}\}$  (14 examples)

$I_3: \{\neg \text{call}, \text{at\_home}\}$  (36 examples)

The first corresponds to the possible worlds where *at\_home* prevents the call regardless of *alarm* (50% of the examples corresponding to the probability of *at\_home* = *false*). The second is the case where both are true and the third where *alarm* is false. The first iteration of the EM algorithm considers the frequency of each example and the probability according to the model with the probability of *alarm* initialized at 0.4. The iteration 0 then associates to  $I_1$  the current probability of *alarm* being either true or false, that is, 0.4. On the other hand, in  $I_2$  (resp.  $I_3$ ) *alarm* has probability 1 (resp. 0) of being true (E step). The updated probability that maximizes the probability of the set of training examples is thus  $\frac{50 \cdot 0.4 + 14 \cdot 1 + 36 \cdot 0}{100} = 0.34$  (M step). Iteration 1 then repeats the E step with the updated probability  $0.34 :: \text{alarm.}$  The probabilities of *alarm* in  $I_1, I_2,$  and  $I_3$  are, respectively, 0.34, 1.0, and 0 (E step) and the new updated probability is  $\frac{50 \cdot 0.34 + 14 \cdot 1 + 36 \cdot 0}{100} = 0.31$  (M step). Further iterations converge to the true probability of  $0.28 :: \text{alarm.}$

*Causal effects.* An alternative view of ProbLog that we exploit in this paper is in terms of CP-logic, a logical language for representing probabilistic causal laws. Probabilistic causal laws model the probability distribution of a set of random variables that are related by a causal process, that is, the variables interact through a sequence of non-deterministic or probabilistic events. ProbLog and CP-logic are closely related as the syntax of a ProbLog program is similar to that of a CP-theory in which each rule has precisely one head atom. Moreover, the semantics of a ground ProbLog program coincides completely with the semantics of CP-Logic. This close relation makes available in ProbLog a modeling technique from CP-logic that we exploit in this paper: the *inhibition effect* (Meert and Vennekens 2014) and *negations in the head* (Vennekens 2013).

Negation in the head gives an interpretation of the negation of heads from ASP in the context of epistemic reasoning and logic theories defining causal mechanisms. In the rest of the paper we adopt this interpretation for negated heads, where the goal is to describe epistemic causal effects and not to obtain classical negation from negation as failure. Vennekens (2013) discusses in detail the differences with classic negation in ASP. In the context of PLP, each rule  $\neg h \leftarrow b_1, \dots, b_n.$  is interpreted by replacing all heads  $h$  in the program with a new atom  $h_{pos}$  and all heads  $\neg h$  with a new atom  $h_{neg}$ , and the rule  $h \leftarrow h_{pos}, \sim h_{neg}.$  is added. This last rule thus defines that  $h$  can be inferred from any of the causes of  $h$  (making  $h_{pos}$  true) only if there is no cause for believing in  $h_{neg}$ , the opposite of  $h$ . The inhibition effect is the resulting decrease in the probability of  $h$  when the probability of  $\neg h$  ( $h_{neg}$ ) increases.

The *noisy-or* effect (Meert and Vennekens 2014) describes the aggregation of the causal probabilities for an atom which appears as the head of multiple (causal) rules.

*Example 7*

We adapt the travel example from Meert and Vennekens (2014) to describe the inhibition effect in ProbLog. Alice is traveling during a pandemic and therefore there is a chance of getting infected. This probability is influenced by two factors: first, if Alice travels with public means of transport there is some probability of being infected there. Second, if Alice is vaccinated then there is a chance that she will not be infected. We encode this last rule with the inhibition effect, which describes a decrease in the probability of infection when vaccinated.

---

0.2 :: *riskyTravel*(alice).  
 0.7 :: *vaccinated*(alice).  
 0.4 :: *infected*(alice)  $\leftarrow$  *riskyTravel*(alice).  
 0.33 ::  $\neg$ *infected*(alice)  $\leftarrow$  *vaccinated*(alice).

---

The two rules correspond to the following rewriting:

---

0.4 :: *infected\_pos*(alice)  $\leftarrow$  *riskyTravel*(alice).  
 0.33 :: *infected\_neg*(alice)  $\leftarrow$  *vaccinated*(alice).  
*infected*(alice)  $\leftarrow$  *infected\_pos*(alice),  $\sim$ *infected\_neg*(alice).

---

*Stable models.* In this paper we will extend ProbLog's semantics, based on two-valued unique well-founded models, with stable model semantics (Gelfond and Lifschitz 1988), where models are two-valued, but a logic program can have more than one stable model. Given a normal logic program  $\mathcal{L}$  and a set  $S$  of atoms interpreted as *true*, called *candidate model*, stable model semantics is defined in terms of the *reduct* of a program  $\mathcal{L}$  w.r.t.  $S$ . The *reduct* of  $\mathcal{L}$  w.r.t.  $S$ ,  $\mathcal{L}^S$ , is defined as follows: (1) for all  $r \in \mathcal{L}$  remove  $r$  from  $\mathcal{L}$  if  $a \in S$  is negated in the body of  $r$ ; (2) for all  $r \in \mathcal{L}$  remove from the body of  $r$  all  $\neg a$  if  $a \notin S$ . If  $S$  is a minimal model for  $\mathcal{L}^S$  then  $S$  is a *stable model* (*answer set*) for  $\mathcal{L}$ .

## 2.2 Probabilistic argumentation

We will consider the application of PLP and SMProbLog to probabilistic argumentation problems (Hunter et al. 2021). An *abstract argumentation framework* (or *argument graph*) (Dung 1995) is a pair  $(A, R)$  where  $A$  is a set of arguments and  $R \subseteq A \times A$  is a binary (attack) relation over  $A$ . Figure 2 represents an example from Hunter et al. (2020) of argument graph  $(A, R)$  where  $A = \{a_1, a_2, a_3, a_4\}$  and  $R = \{(a_1, a_2), (a_2, a_1), (a_3, a_1), (a_4, a_2)\}$ .

In recent years several extensions and modifications of this formulation have been proposed in order to encode more complex argumentation reasoning systems. Among the most relevant, we find *bipolar* argument frameworks (Amgoud et al. 2004), *weighted*




Fig. 2. Abstract argumentation framework. Edges represent attacks, nodes are arguments.




Fig. 3. Probabilistic abstract argumentation framework. Edges represent attacks, nodes are arguments (cfr. Figure 2).

argumentation frameworks (Amgoud et al. 2017), and *probabilistic* argumentation frameworks (Li et al. 2011).

Bipolar argumentation frameworks introduce a *support* relation  $R^+$  along with the attack relation  $R^-$ . A bipolar argument graph is thus a triple  $(A, R^-, R^+)$ , where  $A$  is a set of arguments and  $R^- \subseteq A \times A$  (resp.  $R^+ \subseteq A \times A$ ) is a binary attack (support) relation over  $A$ .

Weighted argumentation frameworks  $(A, R, w)$  augment the traditional argument graph  $(A, R)$  with a weighting function  $w : A \rightarrow [0, 1]$  and belong to a general class of *gradual* argumentation frameworks (Cayrol and Lagasque-Schieux 2005), where arguments' acceptability is evaluated on a fine-grained numerical scale or ranking.

Gradual and bipolar frameworks are combined into *quantitative bipolar argumentation frameworks* (QBAFs) (Baroni et al. 2019), where each argument has a (possibly empty) set of attackers, a (possibly empty) set of supporters, and an initial evaluation (possibly the same for all arguments) on a chosen scale. These elements contribute to a final argument evaluation, provided by a strength function on the chosen scale. A QBAF is thus a quadruple  $(A, R^-, R^+, \tau)$  consisting of a finite set  $A$  of arguments, a binary (attack) relation  $R^-$  on  $A$ , a binary (support) relation  $R^+$  on  $A$ , and a total function  $\tau : A \rightarrow I$ , where  $I$  is the chosen scale.  $\tau(a)$  is thus the base score of  $a$ , for any  $a \in A$ .

Finally, a *probabilistic argument graph* (Li et al. 2011) is a tuple  $(A, R, P_A, P_R)$  where:  $(A, R)$  is an argument graph and  $P_A$  and  $P_R$  are functions:  $P_A : A \rightarrow [0, 1]$  and  $P_R : A \times A \rightarrow [0, 1]$ . Figure 3 describes a probabilistic argument graph derived from the example in Figure 2. The values assigned by the functions  $P_A$  and  $P_R$  are usually interpreted with two different approaches: the *constellations* approach and the *epistemic* approach (Hunter 2013). The constellations approach interprets the probabilities as uncertainty over the structure of the graph and models a probability distribution over subgraphs. Each subgraph is evaluated according to the traditional extension-based semantics (Dung 1995). An extension is a subset of the arguments that is acceptable w.r.t. a given criterion, therefore a distribution over acceptable arguments is derived from (1)

the probability distribution over the possible subgraphs, and (2) the subset of acceptable arguments in each subgraph. The *epistemic* approach on the other hand interprets the probabilities as a direct measure of an agent's belief in the arguments. The accepted arguments are those having probability higher than 0.5, and the argument graph is used to reason about the consistency of the agent's beliefs with respect to the relations encoded in the graph.

### 3 Modeling probabilistic argumentation

Let us now answer the question as to whether PLP can effectively model argumentation processes and reason over their intrinsic uncertainty. Our answer takes the form of encoding argumentation problems and their uncertainty in probabilistic logic programs. By encoding arguments and uncertainty using PLP, we embed argumentation in probabilistic approaches to automated reasoning under uncertainty. More specifically, probabilistic logic programs represent the joint probability distribution. Modeling the joint probability distribution is central to the success of probabilistic graphical models (Koller and Friedman 2011), which are the most popular tools for learning and reasoning about uncertainty in AI. Naively specifying joint probability distribution however requires a number of parameters exponential in the number of random variables (Pearl 1989; Charniak 1991). The success of probabilistic graphical models and Bayesian networks (Pearl 1989) is based on the use of a graphical structure that represents a set of conditional independence assumptions, which allow to compactly encode the joint probability distributions in a factorized form. In Bayesian networks, for instance, the key assumption is that each variable is conditionally independent of its non-descendants, given its parents. Two random variables  $X$  and  $Y$  are said to be *conditionally independent given  $Z$* , a third variable, if in a joint probability distribution  $P$  including the three variables  $P(X|Y, Z) = P(X|Z)$  whenever  $P(Y, Z) > 0$ , that is,  $Y$  does not provide extra information  $X$ , if we already know  $Z$ . In argumentation, conditional probabilities such as  $P(X|Z)$  can be used to update the belief in  $X$  in the light of the argument  $Z$ , while conditional independencies can be used to reason about the (ir)relevance of a certain argument  $Y$  in the context of  $P(X|Z)$ .

By approaching probabilistic argumentation problems from a PLP perspective, we apply the principles in modeling uncertainty in AI formulated by Pearl (1989):

We will also stress that probability theory is unique in its ability to process context-sensitive beliefs, and what makes the processing computationally feasible is that the information needed for specifying context dependencies can be represented by graphs and manipulated by local propagation.

This suggests an interpretation of the argument graph as a probabilistic graphical model, that is, a graph modeling the conditional independencies between random variables (arguments). However, Bayesian networks are directed acyclic graphs, and with argument graphs we are concerned with more complex relations involving cycles and negations (attacks). This motivates the need for a richer framework, PLP, and semantics (Section 4) to derive the joint probability distribution. The argument graph thus expresses how conditional independencies and context dependencies shape a joint belief distribution across

the arguments. We thus follow the epistemic interpretation of probabilities where probabilities measure the degree of belief in the arguments. As Hunter (2013) remarks, the emphasis in the epistemic approach is to find a probability distribution that is rational with respect to the argument graph. With PLP, the derivation of the joint probability distribution from the argument graph mirrors the process of evaluating personal beliefs described by Nilsson (2014):

If we were to examine the relationships among all of our beliefs carefully, an impossible task in practice but one that is interesting to think about, we would see that some of them should make others more credible and some less. They would even compete among themselves with conflicting influences. We can imagine all of the beliefs in our large network of beliefs “fighting it out” to agree finally on the strength of each belief in the network. When they do finally agree, we say that the beliefs *cohere*.

We begin by defining the type of models that describe a probabilistic argumentation problem, that is, we encode the problem in a probabilistic logic program. Following the epistemic interpretation of probabilities we map arguments to logical atoms whose probability (belief) can be queried in the traditional PLP style, that is, querying marginal and conditional probabilities (beliefs) given evidence. Therefore, we do not define an additional layer of logic rules specifying argumentation semantics over a (probabilistic) graph, as in the constellations approach. Mapping a probabilistic argument graph to a probabilistic logic program means defining its semantics, because the PLP semantics of choice for the program in turn defines the graph’s semantics. For this reason, we first discuss the informal semantics of the mapping, and then characterize its formal semantics by means of PLP (Section 4). In the following sections we describe in detail a PLP system that is capable of reasoning over such semantics.

*Mapping.* Given a probabilistic argument graph  $G = (A, R^-, P_A, P_{R^-})$ , we associate each argument  $a \in A$  to a random variable  $arg(a)$  representing the belief in  $a$ . This belief is described by means of the two probability functions:  $P_A$  and  $P_{R^-}$ . We interpret  $P_A(a)$  as the prior belief for  $a$ , that is, a bias independent of the other arguments, similar to a base score in a QBAF. The independence assumption of the prior beliefs in the arguments does not mean that arguments are assumed to be independent, but only the degree to which the agent is biased towards them. The *posterior* belief in an argument is determined by a combination of the bias and the attacks received. We interpret an attack  $P_{R^-}(a, b)$  as the belief that accepting  $a$  causes the rejection of  $b$ . Since we measure beliefs by means of random variables, the natural interpretation of attacks is that the belief in the target is (negatively) conditioned by the belief in the attacker, therefore believing in  $a$  causes a decrease in the belief of  $b$ . Therefore, the belief in an argument depends on the belief in the attacking arguments: a strong belief in any of the attackers can cause the agent to not believe the attacked argument. In Section 2 we showed how in PLP such causal relation is expressed by means of the *inhibition effect*, therefore we model attacks with negation in the heads.

#### Definition 4

Given a probabilistic argument graph  $G = (A, R^-, P_A, P_{R^-})$  the corresponding probabilistic logic program  $\mathcal{L}$  is a program where:

- for each  $a \in A$  and  $P_A(a) = p$ :
  - $p :: bias(a). \in \mathcal{L}$
  - $arg(a) \leftarrow bias(a). \in \mathcal{L}$
- for each  $(a, b) \in R^-$  and  $P_{R^-}((a, b)) = p$ :
  - $p :: \neg arg(b) \leftarrow arg(a). \in \mathcal{L}$

*Example 8*

Consider the example in Figure 3: the corresponding program is:

---

$0.6 :: bias(a_1).$	$0.7 :: bias(a_2).$
$0.4 :: bias(a_3).$	$0.2 :: bias(a_4).$
$arg(a_1) \leftarrow bias(a_1).$	$arg(a_2) \leftarrow bias(a_2).$
$arg(a_3) \leftarrow bias(a_3).$	$arg(a_4) \leftarrow bias(a_4).$
$0.9 :: \neg arg(a_1) \leftarrow arg(a_3).$	$0.8 :: \neg arg(a_2) \leftarrow arg(a_4).$
$0.4 :: \neg arg(a_1) \leftarrow arg(a_2).$	$0.7 :: \neg arg(a_2) \leftarrow arg(a_1).$

---

A rule  $\neg arg(b) \leftarrow arg(a)$  therefore says that believing argument  $a$  causes not believing in  $b$ . In Section 2 we showed how the negation in the head is interpreted in PLP by rewriting the rule as:  $arg_{neg}(b) \leftarrow arg(a).$   $arg(b) \leftarrow arg_{pos}(b), \sim arg_{neg}(b).$  expresses “believing  $a$  causes believing in a counterargument for  $b$  and  $b$  cannot be believed when a counterargument is believed.” The probability  $p = P_{R^-}((a, b))$  measures the belief in the effectiveness of the attack, that is, how likely it is to conclude that argument  $b$  is rejected given that  $a$  is believed.

*Example 9*

The rules in Example 8 are internally rewritten as follows.

---

$0.6 :: bias(a_1).$	$0.7 :: bias(a_2).$
$0.4 :: bias(a_3).$	$0.2 :: bias(a_4).$
$arg_{pos}(a_1) \leftarrow bias(a_1).$	$arg_{pos}(a_2) \leftarrow bias(a_2).$
$arg_{pos}(a_3) \leftarrow bias(a_3).$	$arg_{pos}(a_4) \leftarrow bias(a_4).$
$0.9 :: f_1.$	$0.4 :: f_2$
$0.8 :: f_3$	$0.7 :: f_4$
$arg(a_1) \leftarrow arg_{pos}(a_1), \sim arg_{neg}(a_1).$	$arg_{not}(a_1) \leftarrow arg(a_3), f_1.$
$arg(a_2) \leftarrow arg_{pos}(a_2), \sim arg_{neg}(a_2).$	$arg_{not}(a_1) \leftarrow arg(a_2), f_2.$
$arg(a_3) \leftarrow arg_{pos}(a_3), \sim arg_{neg}(a_3).$	$arg_{not}(a_2) \leftarrow arg(a_4), f_3.$
$arg(a_4) \leftarrow arg_{pos}(a_4), \sim arg_{neg}(a_4).$	$arg_{not}(a_2) \leftarrow arg(a_1), f_4.$

---

We started by considering basic probabilistic argument graphs, but PLP can encode much more complex relations over random variables, hence we exploit the flexibility and expressivity of PLP to effortlessly extend this modeling technique to incorporate many relevant aspects of argumentation in a framework capable of encoding:




Fig. 4. Example 10 representation. Dashed (resp. solid) edges represent supports (resp. attacks)  $R^+$  (resp.  $R^-$ ). Nodes (resp. edges) are labeled with the corresponding bias (belief).

- Quantitative evaluations of relations.
- Support relations.
- Attack (and supports) from sets of arguments.
- Distinctions between proponents of arguments.

We showed in Example 8 how annotated rules express a fine-grained quantification of how strong the causal relation between the arguments (random variables) is.

*Support.* Support relations are symmetric to attacks: we consider the (bipolar) extension of a probabilistic argument graph with a relation support  $R^+$  ( $R^+ \subseteq A \times A, R^+ \cap R^- = \emptyset$ ) and the corresponding probability  $P_{R^+} : R^+ \rightarrow [0, 1]$ . Believing in a supporter  $a$  causes believing in  $b$  with probability  $p$ : this corresponds to an increase of the belief in the supported argument  $b$ . This relation is encoded as a simple inference rule:

*Definition 5*

Given a bipolar probabilistic argument graph  $G = (A, R^-, R^+, P_A, P_{R^-}, P_{R^+})$  the corresponding probabilistic logic program  $\mathcal{L}$  is a program where:

- for each  $a \in A$  and  $P_A(a) = p$ :
  - $p :: bias(a). \in \mathcal{L}$ ,
  - $arg(a) \leftarrow bias(a). \in \mathcal{L}$ ;
- for each  $(a, b) \in R^-$  and  $P_{R^-}((a, b)) = p$ :
  - $p :: \neg arg(b) \leftarrow arg(a). \in \mathcal{L}$ ;
- for each  $(a, b) \in R^+$  and  $P_{R^+}((a, b)) = p$ :
  - $p :: arg(b) \leftarrow arg(a). \in \mathcal{L}$ .

We have thus introduced all the elements required to describe Example 1:

*Example 10*

Consider Example 1: we model  $G = (A, R^-, R^+, P_{R^-}, P_{R^+})$  (see Figure 4) with arguments  $a_i \in A$ , attacks  $R = \{(a_1, a_6), (a_4, a_1)\}, (a_1, a_2), (a_2, a_1)\}$ , and supports  $S = \{(a_5, a_4), (a_3, a_1)\}$ . We also define the following probability functions to derive an example of a probabilistic argumentation problem:  $P_A(a_1) = 0.4, P_A(a_2) = 0.8, P_A(a_3) = 0.3, P_A(a_4) = 0.7, P_A(a_5) = 0.6, P_A(a_6) = 0.7, P_{R^-}(a_1, a_6) = 0.6, P_{R^-}(a_2, a_1) = 0.8, P_{R^-}(a_1, a_2) = 0.7, P_{R^-}(a_4, a_1) = 0.3, P_{R^+}(a_5, a_4) = 0.6, P_{R^+}(a_3, a_1) = 0.5$ .

Which we can thus encode as:

*Example 11*

Example 10 is encoded in the following probabilistic logic program:

---


$$\begin{array}{ll}
 0.4 :: \textit{bias}(a_1). & 0.8 :: \textit{bias}(a_2). \\
 0.3 :: \textit{bias}(a_3). & 0.7 :: \textit{bias}(a_4). \\
 0.6 :: \textit{bias}(a_5). & 0.7 :: \textit{bias}(a_6). \\
 \textit{arg}(A) \leftarrow \textit{bias}(A). \\
 0.6 :: \neg \textit{arg}(a_6) \leftarrow \textit{arg}(a_1). & 0.6 :: \textit{arg}(a_4) \leftarrow \textit{arg}(a_5). \\
 0.3 :: \neg \textit{arg}(a_1) \leftarrow \textit{arg}(a_4). & 0.5 :: \textit{arg}(a_1) \leftarrow \textit{arg}(a_3). \\
 0.8 :: \neg \textit{arg}(a_1) \leftarrow \textit{arg}(a_2). & 0.7 :: \neg \textit{arg}(a_2) \leftarrow \textit{arg}(a_1).
 \end{array}$$


---

*Sets.* In this context, it is immediate to generalize relations between pairs of arguments to relations involving multiple arguments. This is the case of argument systems (Nielsen and Parsons 2006), where the attack between two arguments is generalized to an attack relation from a set of arguments toward a single argument (set-attacks). We can model this in PLP by conjoining attackers (resp. supporters) in the body of rules.

*Example 12*

$\neg \textit{arg}(a_1) \leftarrow \textit{arg}(a_4), \textit{arg}(a_5)$ . is a causal relation effective only when all the arguments in the set of attackers (body) are believed.

Also in this case we can easily combine different argumentative features with each other, for example:

*Example 13*

Examples of gradual joint relations in the setting of Example 1:

- gradual set attacks:  $0.6 :: \neg \textit{arg}(a_1) \leftarrow \textit{arg}(a_4), \textit{arg}(a_5)$ .
- gradual set supports:  $0.3 :: \textit{arg}(a_3) \leftarrow \textit{arg}(a_4), \textit{arg}(a_5)$ .

*PLP tools.* Note that in Example 11 we use first-order predicates to compactly define the relation between arguments and their bias. This is a modeling advantage offered by PLP, whose flexibility can also be exploited to model less standard scenarios. For instance, in the following example we distinguish between different proponents of the arguments and define the bias in terms of a measure of the trust in them:

*Example 14*

Consider Example 11 and the distinction between arguments coming from a proponent ( $\{a_2, a_4, a_5, a_6\}$ ) and those introduced by an opponent ( $\{a_1, a_3\}$ ). We can model bias based on the source of the argument as follows:

---


$$\begin{array}{ll}
 0.4 :: \textit{proponent}. & 0.7 :: \textit{opponent}. \\
 \textit{arg}(a_1) \leftarrow \textit{opponent}. & \textit{arg}(a_2) \leftarrow \textit{proponent}. \\
 \textit{arg}(a_3) \leftarrow \textit{opponent}. & \textit{arg}(a_4) \leftarrow \textit{proponent}. \\
 \textit{arg}(a_5) \leftarrow \textit{proponent}. & \textit{arg}(a_6) \leftarrow \textit{proponent}.
 \end{array}$$


---



First-order logic allows for example expressing compactly that the more agents propose the same argument, the higher its bias is. Assume we express proponents and trust as  $p_1 :: prop(1) \dots p_n :: prop(n)$ . and  $proposes(i, a)$  denotes that argument  $a$  is backed by proponent  $i$ , then a rule  $bias(A) \leftarrow proposes(P, A), prop(P)$ . defines a contribution to the bias of each argument from all its proponents, denoted by each belief  $p_i$ .

We thus showed how PLP offers an expressive language that is capable to encode complex relations between arguments and their beliefs. Ideally, we would like to be able to perform the inference and learning tasks over such programs, and apply the traditional PLP algorithms and tools to models of argumentation problems. This means answering marginal probability queries of the kind: “what is the inferred belief in an argument given its bias and the relationships with the other arguments?” or answering conditional queries by reasoning over evidence, that is, “how does the belief in an argument change if the validity or falsehood of a set of arguments is determined?”. The natural question arising at this point is whether the syntax is paired with semantics and reasoning systems that support this type of models. Unfortunately, the answer is negative: the assumptions that PLP semantics impose on the input program are too restrictive for a broad class of argumentation problems. In particular, the presence of cycles through negation in a probabilistic program  $\mathcal{L} = F \cup R$  usually results in a violation of the requirement of the existence of a single two-valued well-founded model for all the possible (deterministic) logic programs  $\omega \cup R$ ,  $\omega \subseteq F$ . This situation is often present in the case of reciprocal (or, in general, cyclic) attacks in argument graphs. Example 15 shows how arguments attacking each other results in a cyclic relation where negation is involved:

*Example 15*

In Example 11 rules  $0.8 :: \neg arg(a_1) \leftarrow arg(a_2)$ . and  $0.7 :: \neg arg(a_2) \leftarrow arg(a_1)$ . correspond to the following rewriting of the negated heads:

$$\begin{aligned} 0.7 :: f. \quad & arg_{neg}(a_2) \leftarrow f, arg(a_1). \quad arg(a_1) \leftarrow arg_{pos}(a_1), \sim arg_{neg}(a_1). \\ 0.8 :: g. \quad & arg_{neg}(a_1) \leftarrow g, arg(a_2). \quad arg(a_2) \leftarrow arg_{pos}(a_2), \sim arg_{neg}(a_2). \end{aligned}$$

The negation in the head thus leads to a cyclic dependency, we use  $\xrightarrow{\sim}$  to denote a dependency through negation,  $\rightarrow$  otherwise:

$$arg(a_1) \xrightarrow{\sim} arg_{neg}(a_1) \rightarrow arg(a_2) \xrightarrow{\sim} arg_{neg}(a_2) \rightarrow arg(a_1),$$

which causes the program to not satisfy the requirement of a unique two-valued well-founded model in the possible worlds where  $f, g$ , and the probabilistic facts corresponding to the supports of the two arguments are included.

In the Section 4 we discuss why such programs are not captured by traditional PLP semantics, and we propose a new semantics that, on the contrary, allows us to reason over this particular class of programs.

#### 4 Probabilistic semantics

In this section we present a novel PLP semantics to reason over the programs obtained from the mapping of probabilistic argumentation frameworks described in Section 3. In fact, traditional PLP frameworks cannot reason over these programs because they may define cyclic dependencies through negation (Example 15). In PLP, in particular

ProbLog (Fierens *et al.* 2011), the requirement of a single two-valued well-founded model is always satisfied when the program does not define cyclic relations through negation (a sufficient but not necessary condition). In fact, in the case of well-founded semantics for normal logic programs, cycles through negation can result in a unique model where some atoms are labeled with a third value “undefined” (Gelder *et al.* 1991). At the same time, in stable model semantics, which is two-valued, cycles through negation can lead to one, many, or zero stable models (Gelfond and Lifschitz 1988).

*Example 16*

Gelder *et al.* (1991) shows that the program  $a \leftarrow \sim b, b \leftarrow \sim a$ . has no two-valued well-founded model, hence the three-valued well-founded model considers both  $a$  and  $b$  as *undefined*. On the other hand, the program has two stable models, namely  $\{a, \neg b\}$  and  $\{\neg a, b\}$ .

For this reason in ASP non-deterministic choices between atoms can be expressed by recursive dependencies through negation. The answer sets of a program are possible sets of beliefs that a rational agent may hold on the basis of the information expressed by the logic rules (Gelfond and Lifschitz 1991). For this reason, we will consider stable model semantics in order to handle non-probabilistic choices introduced by the logic component.

The case with multiple stable models corresponds to the expression of non-deterministic aspects of the problem modeled by the logic rules (Saccà and Zaniolo 1990).

*Example 17*

Consider Example 2: if we think that Diane will find Hotel X too expensive with probability 0.6 and that Hotel Y will be noisy with probability 0.7, we can write the following program, with the equivalent rewriting on the right:

PROGRAM:	INTERNAL REWRITING:
0.6 :: <i>too_expensive_X</i> .	0.6 :: <i>too_expensive_X</i> .
0.7 :: <i>too_noisy_Y</i> .	0.7 :: <i>too_noisy_Y</i> .
<i>stay_at_Y</i> $\leftarrow$ <i>too_expensive_X</i> .	<i>stay_at_Y_pos</i> $\leftarrow$ <i>too_expensive_X</i> .
<i>stay_at_X</i> $\leftarrow$ <i>too_noisy_Y</i> .	<i>stay_at_X_pos</i> $\leftarrow$ <i>too_noisy_Y</i> .
$\neg$ <i>stay_at_Y</i> $\leftarrow$ <i>stay_at_X</i> .	<i>stay_at_Y_neg</i> $\leftarrow$ <i>stay_at_X</i> .
$\neg$ <i>stay_at_X</i> $\leftarrow$ <i>stay_at_Y</i> .	<i>stay_at_X_neg</i> $\leftarrow$ <i>stay_at_Y</i> .
	<i>stay_at_X</i> $\leftarrow$ <i>stay_at_X_pos</i> , $\sim$ <i>stay_at_X_neg</i>
	<i>stay_at_Y</i> $\leftarrow$ <i>stay_at_Y_pos</i> , $\sim$ <i>stay_at_Y_neg</i>

Intuitively, if the counterarguments for both hotels are considered, that is, the possible world where *too\_expensive\_X* and *too\_noisy\_Y* are included, then we are left with a choice between hotel X and hotel Y. In fact, while the first two rules of the program suggest that Diane should stay at both hotels, the last two rules state that this is logically not possible and only one of the two should be chosen. Therefore, in this possible world logic defines a non-deterministic choice that was not modeled by probabilistic facts.

This case is not compatible with PLP frameworks because the underlying assumption is that uncertainty is fully captured by each total choice and therefore that a user exactly knows the causes and effects of all relevant non-deterministic events that might happen,

for example, “believing in argument  $a$  determines a choice between arguments  $b$  and  $c$ ”. However, we showed in Section 3 (for instance Example 8 and 11) that this is not always the case and obtaining such complete information over the choices may require inference itself. For instance, in argumentation some choices are induced by the relations between arguments (Example 15) and may be required only if some particular conditions apply, for example, not accepting an argument, which may depend in turn on other probabilistic aspects of the problem. Stable model semantics provides a solution to handling non-probabilistic choices introduced by the logic component of a program. At the same time, the absence of a stable model signals that a given selection of probabilistic facts  $F$  (total choice) is inconsistent with the knowledge encoded in the rules  $R$  of the program.

A program with zero stable models reflects the case where it is not possible to define an interpretation of the atoms that satisfies all logic rules, therefore a program that has at least one stable model is called “consistent” (Gelfond 2008).

#### Example 18

Consider a program modeling the famous barber paradox: in a village the barber shaves all men, and those only, who do not shave themselves, does the barber shave himself?

---


$$0.5 :: \text{barber}(\text{bob}). \quad 0.5 :: \text{villager}(\text{bob}).$$

$$\text{shaves}(X, Y) \leftarrow \text{barber}(X), \text{villager}(Y), \sim \text{shaves}(Y, Y).$$


---

In the possible world where  $\text{bob}$  is both a villager and the barber (with probability 0.25), we have an inconsistency: there is neither a two-valued well-founded model nor a stable model. Therefore, we will say that the program is inconsistent with probability 0.25 and the atoms for the inconsistent possible worlds are interpreted with a third value “inconsistent.” This requires us to point out a relevant consideration about the semantics of total choices.

In the past total choices have been defined in two different ways, as an *interpretation* or a *selection* of facts. They are equivalent under the assumption of a one-to-one correspondence between a total choice and the least model or the two-valued well-founded model, but this is not the case when considering a more general semantics for models. In the original definition by Sato (1995) a total choice is an *interpretation* (an assignment of truth values) of the set of probabilistic facts, whose probability distribution is extended to a distribution over least models. Fierens et al. (2011), on the other hand, define in ProbLog an atomic choice ( $p : f$ ) as the choice of *selecting*  $f$  for being included with probability  $p$  or discarded with probability  $1 - p$ . A total choice for a ProbLog program  $\mathcal{L} = F \cup R$  (probabilistic facts  $F$  plus rules  $R$ ) is thus any subset  $\omega \subseteq F$  of the facts, which defines a logic program  $\omega \cup R$ . These two definitions are equivalent under two-valued well-founded semantics because facts are always interpreted as true. This holds also when considering programs with at least one stable model. However, if there is no stable model, it is not possible to consider a total choice an interpretation of the (probabilistic) facts.

#### Example 19

In Example 18 defining the total choice as an interpretation assigning true to both probabilistic facts does not generalize to a normal logic program without (two-valued)

models, because such interpretation is no longer a subset of any (two-valued) model. The total choice  $\{barber(bob), villager(bob)\}$  cannot be regarded as a partial interpretation where both are true because there is no model that extends such interpretation. On the contrary, considering the total choice as a selection of facts does not make any assumption on the eventual interpretation.  $\{barber(bob), villager(bob)\}$  can be thus regarded as an inconsistent choice for the given program and the logical atoms can be interpreted as inconsistent instead of *true*.

Therefore, in this paper we adopt the definition of total choices as subsets of (probabilistic) facts, which thus makes no assumption about their interpretation in the corresponding model(s). More specifically, our semantics uses the following principles to deal with the novel setting.

*Zero models.* Inconsistent programs correspond to a loss of probability mass in the interpretations of probabilistic facts, because such probability mass can no longer be associated to a two-valued model. For instance:

*Example 20*

An alarm goes off when a burglary or earthquake happen. The owner however thinks that the alarm was not properly installed and should not go off because of a valid reason. If the alarm does not activate then the owner is right. We can model this situation with the following (probabilistic) normal logic program:

---

$0.5 :: burglary.$	$0.5 :: earthquake.$
$alarm \leftarrow burglary.$	$alarm \leftarrow earthquake.$
$defective \leftarrow alarm, \sim defective.$	$right \leftarrow \sim alarm.$

---

There are four total choices:  $\omega_1 = \{burglary, earthquake\}$ ,  $\omega_2 = \{burglary\}$ ,  $\omega_3 = \{earthquake\}$ ,  $\omega_4 = \{\}$ . The logic part of the program is inconsistent when *alarm* is true, therefore when either *burglary* or *earthquake* are true. This means that for three total choices,  $\omega_1, \omega_2$ , and  $\omega_3$ , the corresponding possible world has neither a two-valued well-founded model nor stable models. In these cases the program is inconsistent with probability  $P(\omega_1) + P(\omega_2) + P(\omega_3) = 0.75$  and the only total choice corresponding to a consistent possible world is  $\omega_4$  which corresponds to the stable model  $\{right\}$ .

We thus assign the probability mass associated to inconsistencies to an “inconsistent” state, whose probability thus quantifies how likely it is to consider an inconsistent possible world among those defined by the program. We define this in terms of a model where all atoms are labeled with a third value “inconsistent.” This solution has the advantage of preserving the probability mass of total choices by assigning it to the model representing an inconsistent possible world. The semantics is thus informative as to the degree of inconsistency of the logical part of the program. Note that this is different from three-valued well-founded semantics, where the atoms involved in the cycles responsible for multiple stable models are labeled with the third value (“undefined”) as well, and not all atoms are interpreted undefined at the same time.

*Example 21*

When arguing about getting vaccinated or not there are different media we can choose to believe in: virologists, newspapers, and Facebook. Virologists endorse vaccines and their safety, newspapers argue that there is some degree of risk involved, Facebook posts reject the acceptability of such degree of risk. We can model this reasoning pattern and the level of trust in the different media with the following program:

---

$0.9 :: \text{trust\_virologists.}$      $0.7 :: \text{trust\_newspapers.}$      $0.3 :: \text{trust\_facebook.}$   
 $\text{safe} \leftarrow \text{trust\_virologists,} \sim \text{dangerous.}$   
 $\text{reasonable\_risk} \leftarrow \text{trust\_newspapers,} \sim \text{safe.}$   
 $\text{dangerous} \leftarrow \text{trust\_facebook,} \sim \text{reasonable\_risk.}$

---

The three rules determine a cyclic dependency (of odd length) with negation involved between *safe*, *reasonable\_risk*, and *dangerous* and therefore trusting all media is an inconsistency, because there is no possible truth assignment to *safe*, *reasonable\_risk*, and *dangerous*. We will then say that this program is inconsistent with probability  $0.9 \cdot 0.7 \cdot 0.3 = 0.189$ . The remaining probability mass of  $(1 - 0.189) = 0.811$  is divided over the other possible words where the total choice of probabilistic facts is a consistent choice according to the logic of the program. These choices thus induce a two-valued assignment to the atoms of the program.

*At least one model.* In order to give a probabilistic interpretation of logical choices, we follow the principle of maximum entropy. The maximum-entropy principle is a widely adopted principle in Bayesian reasoning and statistical relational learning. It states that the probability distribution which best represents the current state of knowledge about a system is the one with the largest entropy (Jaynes 1988). In this setting this principle translates to considering a uniform distribution for the (multiple) stable models of a total choice.

*Example 22*

Consider the following (probabilistic) normal logic program:

---

$0.5 :: \text{burglary.}$                        $0.5 :: \text{earthquake.}$   
 $\text{alarm} \leftarrow \text{burglary.}$              $\text{defective} \leftarrow \text{earthquake.}$   
 $\text{alarm} \leftarrow \sim \text{defective.}$          $\text{defective} \leftarrow \sim \text{alarm.}$

---

There are four total choices:  $\omega_1 = \{\text{burglary, earthquake}\}$ ,  $\omega_2 = \{\text{burglary}\}$ ,  $\omega_3 = \{\text{earthquake}\}$ ,  $\omega_4 = \{\}$ . While the first three correspond to one stable model, that is,  $MOD(\omega_1) = \{\{\text{burglary, earthquake, alarm, defective}\}\}$ ,  $MOD(\omega_2) = \{\{\text{burglary, alarm}\}\}$ ,  $MOD(\omega_3) = \{\{\text{earthquake, defective}\}\}$ ,  $\omega_4$  has two stable models:  $MOD(\omega_4) = \{\{\text{alarm}\}, \{\text{defective}\}\}$ . Since this choice is not related to beliefs but rather to logical consistency, we assume that all stable models are equally probable for the given total choice  $\omega_4$ . In this case we will choose *alarm* half of the times and *defective* the other

half. Since  $P(\omega_4) = 0.25$  the probability of the model  $\{alarm\}$  is equal to the probability of the model  $\{defective\}$ , namely 0.125.

*Example 23*

The program of Example 17 has four consistent possible worlds, one of which with two stable models:

Possible world's choices	$P(\omega)$	Model $M_\omega$	$\hat{P}(M_\omega)$
$\omega_1 = \{\}$	0.12	$M_{\omega_1} = \omega_1$	0.12
$\omega_2 = \{too\_expensive\_X\}$	0.18	$M_{\omega_2} = \omega_2 \cup \{stay\_at\_Y\}$	0.18
$\omega_3 = \{too\_noisy\_Y\}$	0.28	$M_{\omega_3} = \omega_3 \cup \{stay\_at\_X\}$	<b>0.28</b>
$\omega_4 = \{too\_expensive\_X, too\_noisy\_Y\}$	0.42	$M_{\omega_4}^1 = \omega_4 \cup \{stay\_at\_Y\}$	0.21
		$M_{\omega_4}^2 = \omega_4 \cup \{stay\_at\_X\}$	<b>0.21</b>

Consider  $\omega_4$ : when both probabilistic facts *too\_expensive\_X* and *too\_noisy\_Y* are conjoined with the rules, there is no probabilistic justification for choosing one hotel over the other, but rules impose a choice. In this case, we consider the probability of *stay\_at\_X* and *stay\_at\_Y* to be equivalent, because in a possible world where both counterarguments are true, a choice of a hotel is still necessary. The probabilistic part of the program which is quantified by the user is fixed in the total choice  $\omega_4$ , therefore in determining the models of  $\omega_4$  there is no probabilistic justification for assigning one model a probability higher than another. In the definition of the probability distribution of the models  $\hat{P}$ , we thus divide equally the probability of the possible world determined by  $\omega_4$  between its two models:  $\hat{P}(M_{\omega_4}^1) = \hat{P}(M_{\omega_4}^2) = \frac{0.42}{2} = 0.21$ . The probability of an atom *a* remains the sum of the probability of the models where *a* is true, for example the probability of *stay\_at\_X* is  $\hat{P}(M_{\omega_3}) + \hat{P}(M_{\omega_4}^2) = 0.28 + 0.21 = 0.49$ .

**4.1 Formal semantics**

We define the models for each total choice and the corresponding probability in order to extend a probability distribution over total choices to a probability distribution over models. Each total choice  $\omega$  corresponds to: (1) a probability  $P(\omega)$  as defined by distribution semantics and (2) a set  $SM(\omega)$  of stable models, over which the probability gets distributed if it is non-empty. For this distribution to be well-defined, we only consider programs with finite Herbrand base, and thus a finite number of interpretations.

*Definition 6*

A valid SMProbLog program is a probabilistic normal logic program without function symbols.

*Definition 7*

Given a valid SMProbLog program  $\mathcal{L}$ , the probability  $P(\omega)$  of a total choice  $\omega$  is:

$$P(\omega) = \prod_{(f:p) \in \mathcal{L}, f \in \omega} p \cdot \prod_{(f:p) \in \mathcal{L}, f \notin \omega} 1 - p.$$

that is, the product of the probabilities of the facts for being included/excluded from  $\omega$ .

*Example 24*

The probability of the total choices in Examples 20 and 22 is  $P(\omega_1) = P(\omega_2) = P(\omega_3) = P(\omega_4) = 0.5 \cdot 0.5 = 0.25$

*Definition 8*

Given a probabilistic normal logic program  $\mathcal{L} = F \cup R$  and a corresponding total choice  $\omega \subseteq F$ ,  $SM(\omega)$  is the set of stable models of the (deterministic) program  $\omega \cup R$ .

We distinguish two cases:  $SM(\omega) = \emptyset$  and  $|SM(\omega)| > 0$ , and we define accordingly a set of models for a possible world.

A model is a three-valued interpretation  $(T, F)$  where  $T$  is the set of true atoms and  $F$  is the set of false atoms. The remaining atoms are interpreted as being part of an inconsistent possible world.

*Definition 9*

Given a probabilistic logic program  $\mathcal{L}$  and let  $HB(\mathcal{L})$  be its Herbrand base. For each total choice  $\omega \in \Omega_{\mathcal{L}}$ , the corresponding set of models  $MOD(\mathcal{L}, \omega)$  is:

$$MOD(\mathcal{L}, \omega) = \begin{cases} \{(\emptyset, \emptyset)\} & \text{if } |SM(\omega)| = 0. \\ \{(M_\omega, HB(\mathcal{L}) \setminus M_\omega) \mid M_\omega \in SM(\omega)\} & \text{if } |SM(\omega)| > 0. \end{cases}$$

Each total choice thus has at least one model: either the “inconsistent” model or one or more stable models. We extend the probability of the total choices to a distribution over the corresponding models by applying the maximum-entropy principle.

The probability of a model  $M_\omega \in MOD(\mathcal{L}, \omega)$  is the probability of the corresponding total choice  $\omega$  normalized w.r.t. the number of the models for that possible world:

*Definition 10*

Given a probabilistic normal logic program  $\mathcal{L}$  and a total choice  $\omega$ ,  $\forall M_\omega \in MOD(\mathcal{L}, \omega)$ :

$$\hat{P}(M_\omega) = \frac{P(\omega)}{|MOD(\mathcal{L}, \omega)|}.$$

We thus derive a distribution over models  $\hat{P}$  from the probability distribution over total choices  $P$ . We can now define the probability of the program of being inconsistent and of a query of an atom. The degree of inconsistency of the program  $\mathcal{L}$ ,  $\mathbb{P}(\mathcal{L} \vdash \perp)$  is the sum of the probabilities of interpreting the atoms as inconsistent in each possible world:

$$\mathbb{P}(\mathcal{L} \vdash \perp) = \sum_{\omega \in \Omega_{\mathcal{L}}, MOD(\mathcal{L}, \omega) = \{(\emptyset, \emptyset)\}} \hat{P}((\emptyset, \emptyset)).$$

As for the consistent possible worlds, we can query the program for the probability of an atom being interpreted as true (or false):

*Definition 11*

Given a probabilistic normal logic program  $\mathcal{L}$ , the probability of success of querying  $a$  is the probability of interpreting  $a$  true in a consistent possible world:

$$\mathbb{P}(a) = \sum_{a \in M_\omega, (M_\omega, HB(\mathcal{L}) \setminus M_\omega) \in MOD(\mathcal{L}, \omega), \omega \in \Omega_{\mathcal{L}}} \hat{P}(M_\omega).$$

Note that we can derive the probability of an atom being false from the probability of the inconsistent possible worlds and the probability of success:

$$\mathbb{P}(\neg a) = 1 - \mathbb{P}(\mathcal{L} \vdash \perp) - \mathbb{P}(a),$$

because the probability mass of the non-inconsistent models is divided over the two-valued stable models, therefore if an atom is in a consistent possible world, then it is interpreted either true or false.

#### Example 25

We complete Example 20: the probability of the program being inconsistent is  $P(\omega_1) + P(\omega_2) + P(\omega_3) = 0.75$ . The only consistent possible world is  $\omega_4$ , whose only model is  $\{e\}$ :  $\hat{P}(\text{right}) = 0.25$ , therefore  $\mathbb{P}(\text{right}) = 0.25$  and  $\mathbb{P}(\neg \text{right}) = 0$ , vice versa,  $\mathbb{P}(\text{burglary}) = \mathbb{P}(\text{earthquake}) = \mathbb{P}(\text{alarm}) = \mathbb{P}(\text{defective}) = 0$  and  $\mathbb{P}(\neg a) = \mathbb{P}(\neg b) = \mathbb{P}(\neg c) = \mathbb{P}(\neg d) = 0.25$ .

#### Example 26

We complete Example 22: models  $\{\text{burglary}, \text{earthquake}, \text{alarm}, \text{defective}\}$ ,  $\{\text{burglary}, \text{alarm}\}$ ,  $\{\text{earthquake}, \text{defective}\}$  are the unique model for, respectively,  $\omega_1, \omega_2, \omega_3$  hence  $\hat{P}(\{\text{burglary}, \text{earthquake}, \text{alarm}, \text{defective}\}) = \hat{P}(\{\text{burglary}, \text{alarm}\}) = \hat{P}(\{\text{earthquake}, \text{defective}\}) = P(\omega_1) = P(\omega_2) = P(\omega_3) = 0.25$ . The likelihood of  $\omega_4$  is uniformly distributed over the two stable models  $\{\text{alarm}\}$  and  $\{\text{defective}\}$ , whose probability is thus  $\frac{P(\omega_4)}{|\text{MOD}(\mathcal{L}, \omega_4)|} = \frac{0.25}{2} = 0.125$ . Therefore,  $\mathbb{P}(\text{burglary}) = \mathbb{P}(\text{earthquake}) = 0.5$ ,  $\mathbb{P}(\text{alarm}) = \mathbb{P}(\text{defective}) = 0.625$ .

## 5 The joint distribution of beliefs

By defining the semantics of SMProbLog programs we also define the semantics of the mapping from probabilistic argumentation problems to PLP (Section 3). Under these semantics a probabilistic argument graph thus represents a *joint probability distribution* of the arguments' beliefs by means of the corresponding probabilistic logic program. As we argued in Section 3, the joint probability distribution represents the marginal and conditional strengths of the beliefs in the arguments from their biases and relations. That is, we are concerned with determining the probabilities of logically related propositions.

#### Example 27

In Example 10 the marginal probabilities of each argument in the inferred joint probability distribution are:  $\mathbb{P}(\text{arg}(a_1)) = 0.29$ ,  $\mathbb{P}(\text{arg}(a_2)) = 0.63$ ,  $\mathbb{P}(\text{arg}(a_3)) = 0.3$ ,  $\mathbb{P}(\text{arg}(a_4)) = 0.81$ ,  $\mathbb{P}(\text{arg}(a_5)) = 0.6$ ,  $\mathbb{P}(\text{arg}(a_6)) = 0.58$ .

These marginal distributions are thus coherent with the initial biases and relations in the sense of Nilsson (2014). Polberg *et al.* (2017) define several properties for epistemic argument graphs, among which coherency, to evaluate arbitrary assignments of values in  $[0, 1]$  under the principle that a probability higher than 0.5 defines an argument as accepted. In the epistemic approach literature, acceptance is thus a function of the arbitrary probability values with respect to an arbitrary threshold. Marginal probabilities, representing the agent's beliefs, are thus directly assigned to arguments, and they are evaluated with respect to local properties. Hunter *et al.* (2018) generalize this



approach with *epistemic graphs*, where, instead of considering a single assignment, a set of belief assignments is associated to the argument graph such that they satisfy a set of constraints. This does not narrow down the belief distribution to a single joint probability distribution, but admits a set of possible assignments compatible with the constraints.

Therefore, past work evaluates arbitrary belief assignments with respect to an arbitrary threshold: considering a belief *absolute* value of 0.5 the discriminant between accepted and rejected arguments. This poses a well-known question in the PLP literature, “Where do the numbers come from?” (Charniak 1991), in regard of evaluating subjective beliefs with a fixed value of reference. The main difference with our approach is that in our framework acceptance is a function of logic and its (joint) probability is the associated (joint) belief measure. The joint belief (probability) distribution represents how the prior beliefs (biases) interact with each other (arguments’ logical relations) to define the *unique* joint belief distribution that is coherent with the given argumentative structure (argument graph). The joint belief distribution defines how the *relative* strength of the arguments changes when their logical relations are considered, regardless of an acceptance threshold. This expresses two principles of modeling uncertainty in AI. First, as de Finetti (1931) remarks in his discussion on subjective beliefs, rather than evaluating the absolute degree of beliefs, an agent is coherent when the relative differences between such degrees do not contradict probability axioms. Second, we apply the approach of probabilistic graphical models to probability and reasoning: “Probability is not really about numbers; it is about the structure of reasoning” (Glenn Shafer, cited in Pearl (1989)).

This method allows us to learn probabilities (beliefs) from data (Sections 6.2 and 7) and thus ground their nature in the observation of accepted or rejected arguments. Moreover, to study the properties of the joint probability distribution of arguments we can rely on the traditional inference tasks of PLP frameworks. Interestingly, Nilsson (2014) remarks that considering the consequences and the explanations for beliefs is fundamental for critical thinking, and we can do precisely this by means of two typical PLP tasks: conditional queries and most probable explanation (MPE) queries.

*Conditionals.* We can perform conditional reasoning on the joint probability distribution by querying the probability of atoms given some evidence. We can thus analyze the consequences of beliefs by answering conditional queries such as: “What are the consequences of accepting  $a_1$ ?”. Note that  $a_6$  is *conditionally independent* of the other arguments *given*  $a_1$ , because when  $a_1$  is known to be accepted (or rejected) the belief in  $a_6$  depends only on the corresponding attack and bias.

#### Example 28

By querying the model from Example 15 with the addition of the evidence  $arg(a_1) = true$ , usually simply denoted with  $arg(a_1)$ , we infer the following conditional probabilities:  $\mathbb{P}(arg(a_1)|arg(a_1)) = 1$ ,  $\mathbb{P}(arg(a_2)|arg(a_1)) = 0.08$ ,  $\mathbb{P}(arg(a_3)|arg(a_1)) = 0.36$ ,  $\mathbb{P}(arg(a_4)|arg(a_1)) = 0.75$ ,  $\mathbb{P}(arg(a_5)|arg(a_1)) = 0.58$ ,  $\mathbb{P}(arg(a_6)|arg(a_1)) = 0.28$ . We can also condition the distribution on the relations, for instance by adding the evidence that accepting  $arg(a_1)$  always excludes  $arg(a_6)$ . In this case, if we make explicit the belief in the attack with a probabilistic fact *att* we obtain the same distribution except for  $\mathbb{P}(arg(a_6)|arg(a_1), att) = 0$ .

The conditional independence assumption materializes also when we introduce a belief update to an argument which is conditionally independent of another, given the evidence for an argument that represents the connection of the two beliefs.

*Example 29*

Consider again Example 28 with evidence  $\text{arg}(a_1) = \text{true}$ . If we add the evidence  $\text{arg}(a_5) = \text{false}$ , let  $ev = (\text{arg}(a_1) = \text{true} \wedge \text{arg}(a_5) = \text{false})$ , then we obtain the following marginals:  $\mathbb{P}(\text{arg}(a_1)|ev) = 1$ ,  $\mathbb{P}(\text{arg}(a_2)|ev) = 0.08$ ,  $\mathbb{P}(\text{arg}(a_3)|ev) = 0.36$ ,  $\mathbb{P}(\text{arg}(a_4)|ev) = 0.62$ ,  $\mathbb{P}(\text{arg}(a_5)|ev) = 0$ ,  $\mathbb{P}(\text{arg}(a_6)|ev) = 0.28$ . The only beliefs affected by the new information about  $a_5$  are those for  $a_5$  and  $a_4$ , because  $a_3$  is independent of  $a_5$ ,  $a_1$  is known, and  $a_6$  and  $a_2$  are *conditionally independent given*  $a_1$ .

Conditional independence can also be exploited in argumentation to select arguments that provide new information with respect to an argument or claim. If these arguments' uncertainty can be eliminated by means of experiments or observations, the conditional independence assumption allows us to select only those strictly necessary to influence the belief in the argument of interest.

*Example 30*

Bayesian statistical methods are suited for designing medical clinical trials to find an optimal design which minimizes the empirical results required and at the same time maximizes the information provided regarding the quality of the trial (Berry 2006). Given an argumentative graph with respect to some treatment hypothesis, if some experimental results are unknown then they are associated with a probability (degree of belief). The uncertainty can be removed by collecting empirical evidence, but this comes at a cost of time and resources. In this case, the conditional independence of the claim with respect to some experiments given other empirical results can be exploited to select only those experiments that truly affect the belief in the claim.

In SMProbLog it is also possible to condition the theory to be consistent, by providing the evidence that an atom is not interpreted as inconsistent. This allows us to reason about how beliefs change if they are conditioned to be consistent with the logical structure.

*Example 31*

Consider Example 18: if we add the evidence  $\text{villager}(\text{bob}) = \neg\text{inconsistent}$  then we consider only the consistent possible worlds. Therefore, the probability of each possible world is normalized w.r.t. the probability mass of the consistent total choices. Hence, the probability that bob is both a villager and a barber becomes 0, because there are no models for consistent worlds where this is true. Each consistent possible world then increases its probability from 0.25 to  $\frac{1}{3}$ . For instance  $\mathbb{P}(\text{villager}(\text{bob})|\text{villager}(\text{bob}) = \neg\text{inconsistent}) = \frac{1}{3}$ .

*Most probable explanation.* MPE is a typical PLP task that allows us to address the other point of critical thinking, reasoning on explanations for beliefs. The MPE task is to compute the most likely possible world where some given evidence holds.

*Example 32*

If we compute a MPE query for  $\text{arg}(a_1) = \text{true}$ , we are asking which set of beliefs compatible with accepting  $a_1$  is the most probable. This means finding the most likely set

of probabilistic facts representing a choice of biases and relations with the consequence of accepting  $a_1$ . In the MPE world the biases believed to support the corresponding argument are  $\{bias(a_1).bias(a_2).bias(a_4).bias(a_5).bias(a_6).\}$ . All relations are believed to hold, except for the attack of  $a_4$  to  $a_1$ .  $a_1$  is thus independent of  $a_4$  and  $a_5$  and its acceptance is supported by its bias and the non-deterministic choice with  $a_2$  (following the evidence). These choices, that is, the MPE, have probability 0.0093. Their model is  $\{arg(a_1), \neg arg(a_2), \neg arg(a_3), arg(a_4), arg(a_5), \neg arg(a_6)\}$  which is thus compatible with the evidence of accepting  $arg(a_1)$ . The two arguments attacked by  $a_1$ ,  $a_2$ , and  $a_6$  are rejected, along with  $a_3$  because of the choice of not considering a bias toward it.

### 5.1 Epistemic properties

The joint probability distribution describes the subjective assignment of beliefs to arguments coherent with all possible belief choices of prior bias and relations, regardless of arbitrary thresholds. This means that the joint beliefs are globally coherent with the argumentative structure and therefore the marginal beliefs are also locally coherent with the parents and children's beliefs. The properties for epistemic probabilistic argumentation by Polberg et al. (2017) can thus be used to evaluate the bias of the agent, which is indeed an arbitrary assignment of values, but in evaluating the joint probability distribution, having probability higher than 0.5 has no special meaning. We clarify this by considering some of these properties and show that the marginal belief of the arguments defined by the joint distribution is compatible with the given argumentative structure even if they may not satisfy such properties.

*Coherency.* Polberg et al. (2017) define coherency in epistemic probabilistic argumentation as the property that for all attacks  $(a, b)$  believed with probability  $> 0.5$ ,  $P_A(a) \leq 1 - P_A(b)$ . In the joint probability distribution the marginal probability of arguments  $\mathbb{P}(arg(-))$  is coherent not just locally, considering pairwise marginals, but globally, where the structure of all relations and biases is taken into account.

#### Example 33

Consider the case where two arguments  $a$  and  $b$  are such that  $a$  attacks  $b$  with probability 0.6 and both arguments have bias (prior) 1.  $a$  is not attacked so  $P_A(a) = \mathbb{P}(arg(a)) = 1$ . We can say that the initial assignment, the priors, is not coherent in the sense of Polberg et al. (2017) because  $P_A(a) = 1 \not\leq 1 - 1$ . However, the joint probability distribution defines posterior marginals that are coherent in the sense of Nilsson (2014):  $\mathbb{P}(arg(b)) = 0.4$  because the uncertainty about the attack leaves room for possible worlds where the attack is not believed to hold (probability  $1 - 0.6 = 0.4$ ) and thus it is possible to accept both  $a$  and  $b$ . Here  $1 = \mathbb{P}(arg(a)) \not\leq 1 - \mathbb{P}(arg(b)) = 0.6$ , but  $\mathbb{P}(arg(a))$  should not be less or equal than 0.6 because it is independent of  $b$ , while accepting  $b$  is influenced by accepting  $a$  and its attack, and thus its prior bias, not the attacker, is inhibited to 0.6.

*Rational.* A rational (resp. strict) assignment resembles the conflict-free principle that if an attacker  $a$  is accepted than the target  $b$  should not:  $P_R((a, b)) > 0.5$  and  $P_A(a) > 0.5$  implies  $P_A(b) \leq 0.5$  (resp.  $P_A(b) < 0.5$ ). In our framework, when choosing a set of prior beliefs (probabilistic facts) the conflict-freeness of the accepted (true) arguments is always guaranteed by the logic rules. This does not exclude the possibility that both  $a$  and  $b$  in

the joint probability distribution are believed with probability  $> 0.5$ . In fact, when an argument has high bias, for example, we think it is almost certain, only a combination of strong attackers with strong relations can lower its inferred probability under 0.5.

*Example 34*

For example, a prior probability of 0.6 of believing in  $a$  and 0.6 in the attack, yields an inhibition by a factor of 0.36, which is not guaranteed to lower the marginal probability of  $b$  under 0.5. In 36% of the possible worlds  $b$  is rejected because of the presence of  $a$  and the attack, but if  $P_A(b) = 1$  then  $\mathbb{P}(arg(b)) = 0.64$ , together with  $\mathbb{P}(arg(a)) = 0.6$  and  $P_R((a, b)) = 0.6$ .  $\mathbb{P}(arg(b)) \leq 0.5$  only when  $\mathbb{P}(a) \cdot P_R((a, b)) \geq 0.5$  (in the case of inhibition from a single attack). This does not mean that the joint probability distribution is not rational in the sense that it defines a set of accepted arguments not conflict-free. The joint probability distribution reflects the relative strength of the arguments and simply defines a (coherent) probability  $\geq 0.5$  to be in a possible world where either  $a$  or  $b$  are accepted, but there is no possible world where the attack is believed and both are accepted.

*Protective.* A (restricted) protective assignment is such that  $P_R((a, b)) \geq 0.5$  (resp.  $P_R((a, b)) > 0.5$ ) and  $P_A(b) > 0.5$  implies  $P_A(a) \leq 0.5$ . Similarly to the rational property, in epistemic probabilistic argumentation this is the principle of rejecting the attacker if the target is accepted. Again, this principle in the joint belief distribution is not enforced by means of the probability (belief) values, but by means of the logical structure of the attack relation.

*Example 35*

Our previous example also describes a belief distribution that is not protective: while  $\mathbb{P}(b) = 0.64$  and  $P_R((a, b)) = 0.6$ ,  $\mathbb{P}(a)$  remains equal to its original prior  $P_A(a) = 0.6$ , thus  $\mathbb{P}(arg(a)) > 0.5$ . This does not mean that both arguments are accepted *because* their probability is higher than 0.5, but it means that a joint belief induced by the argumentative structure coherent with the prior bias assigns a subjective marginal belief  $\mathbb{P}(arg(a))$  and  $\mathbb{P}(arg(b))$  higher than 0.5 to the arguments. The difference in strength is still coherent with the respective biases and belief in the attack relation and probability axioms.

Similar arguments can be constructed for the other properties from Polberg *et al.* (2017): past epistemic argumentation approaches consider the probabilities of the argument graph as the *marginal* beliefs of the agent and check if those are (locally) rational, coherent, ... Our approach uses this arbitrary assignment as the *prior* set of beliefs, and the argumentative structure, to infer a joint belief distribution that is compatible with the initial arbitrary assignment according to the probability postulates and the logical relations of the arguments. The joint belief (probability) distribution thus represents how the prior beliefs (biases) interact with each other to define *marginal and conditional* beliefs that are coherent with the given priors and argumentative (logic) structure.

## 6 Inference and learning

In this section we describe the inference and learning algorithms for SMProbLog semantics. We consider the inference and learning algorithms implemented in ProbLog2 and




Fig. 5. ProbLog2 inference schema.

adapt them to the semantics defined in Section 4. This allows us to study the advantages and shortcomings of traditional PLP inference and learning techniques when a more general semantics is concerned, and indirectly examine their application to argumentation problems.

### 6.1 Inference task

The state-of-the-art inference technique for probabilistic logic programs is reducing the inference task to a WMC problem (cfr. Section 2), therefore we consider its applicability in the context of our semantics. In ProbLog2 inference is implemented by means of a transformation of the ground program to conjunctive normal form (CNF), which is the input for a knowledge compiler (under well-founded semantics) that returns the logical circuit (e.g., a d-DNNF) such that on the corresponding arithmetic circuit the WMC task can be solved in polynomial time (Figure 5). In the rest of the section we describe the adaptations required to this inference technique required for SMProbLog's semantics.

The normalization of the weight of the models entails considering a variant of the *WMC* problem where models are weighted with the corresponding normalization constant. We denote such constant with  $\hat{w}(M_\omega)$ , where  $M_\omega$  is a model for a total choice  $\omega$ , hence  $\hat{w}(M_\omega) = \frac{1}{|\text{MOD}(\mathcal{L}, \omega)|}$ . The *WMC* problem ( $\text{WMC}_{\mathcal{L}}(\varphi) = \sum_{M \in \text{MOD}(\mathcal{L}), M \models \varphi} \prod_{l \in M} w(l)$ ) thus becomes:

$$\widehat{\text{WMC}}_{\mathcal{L}}(\varphi) = \sum_{M \in \text{MOD}(\mathcal{L}), M \models \varphi} \hat{w}(M) \cdot \prod_{l \in M} w(l).$$

This formulation of the problem requires us to determine the weights  $\hat{w}(M_\omega)$ , which, contrary to the weights of the literals, are not known from  $\mathcal{L}$ . In order to determine the weighting function  $\hat{w}$  we thus need to solve an additional counting problem, namely the task of counting how many models agree on a subset of atoms in  $\mathcal{L}$ , that is, the probabilistic facts. We now describe the reduction pipeline in SMProbLog to a  $\widehat{\text{WMC}}$  problem. The reduction can be broadly divided into four components, represented in Figure 6: (1) grounding the program  $\mathcal{L}$ , (2) compiling the logic theory  $\mathcal{L}$  into a d-DNNF, (3) deriving a circuit for retrieving the number of models corresponding to a total choice and an arithmetic circuit for the  $\widehat{\text{WMC}}_{\mathcal{L}}$  problem, and (4) the evaluation of the two circuits to define the solution of the  $\widehat{\text{WMC}}_{\mathcal{L}}$  problem (Algorithm 2).

(1) *Grounding*. In SMProbLog we adopt a standard bottom-up grounding technique. In ProbLog2, the detection of cycles through negations during the grounding step is used as a sufficient condition to reject invalid programs. On the contrary, we use the absence of cycles through negations as a condition to reduce to ProbLog's evaluation algorithm,

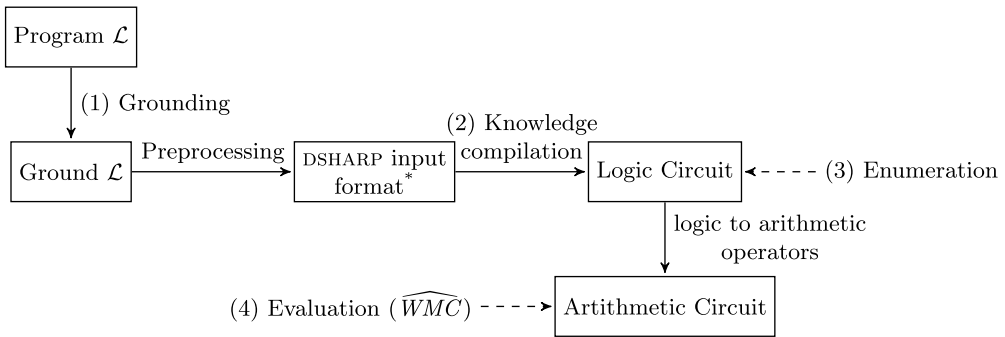


Fig. 6. SMProbLog inference schema. (\*) denotes a different version of DSHARP from ProbLog2, specific for stable model counting.

where SMProbLog's semantics is equivalent, as discussed in Section 8.1. In fact, when this is the case, we skip the enumeration step (Step 3) because  $\hat{w}(M) = 1$  for all stable (and well-founded) models  $M$ . Therefore, the  $\widehat{WMC}_{\mathcal{L}}$  problem becomes equivalent to a  $WMC_{\mathcal{L}}$  problem.

(2) *Compilation.* The goal of the compilation step is to transform the ground propositional theory in a representation that allows us to answer efficiently queries about the probability of atoms. Knowledge compilation produces a logical circuit describing the possible worlds' models of the input propositional theory. Therefore, this step regards exclusively the logical part of the program, while the following inference steps concern the computation of probabilities. As in ProbLog2, the knowledge compiler is a black box with respect to the probabilistic inference system. Therefore, in SMProbLog we are only concerned with transforming the propositional ground theory to the suitable input format. The input format and the implementation details of the knowledge compiler in SMProbLog can be found in Aziz et al. (2015), which extends the DSHARP compiler (Muise et al. 2012) with stable model semantics. In ProbLog2 the input format is a theory in CNF, Aziz et al. (2015) replace it with an encoding in CNF format of the rules, variables, and the corresponding strongly connected components. The output is a logical circuit, a d-DNNF, describing the *stable* models of the input propositional theory. The circuit therefore expresses only the models corresponding to the consistent total choices. This representation, with a standard transformation in an arithmetic circuit, allows efficient stable model counting, that is, solving the  $WMC$  problem. Figure 7 shows the output of the knowledge compilers from the representations of Example 20 and 22. Contrary to ProbLog2, however, the transformation of the logical circuit into an arithmetic circuit is not sufficient to solve the  $\widehat{WMC}$  problem. This because the arithmetic circuit allows us to efficiently compute the solution to a  $WMC$  problem, but the normalization constants  $\hat{w}$  are still unknown. Unless the weighting is determined by the absence of cycles through negations, as defined in the previous paragraph, the enumeration step is required to associate each model to the corresponding normalization constant.

(3) *Enumeration.* The purpose of the enumeration step is to count the number of stable models corresponding to each (consistent) total choice, which defines  $\hat{w}$ . Unfortunately, this operation cannot be performed efficiently on the representation obtained from the

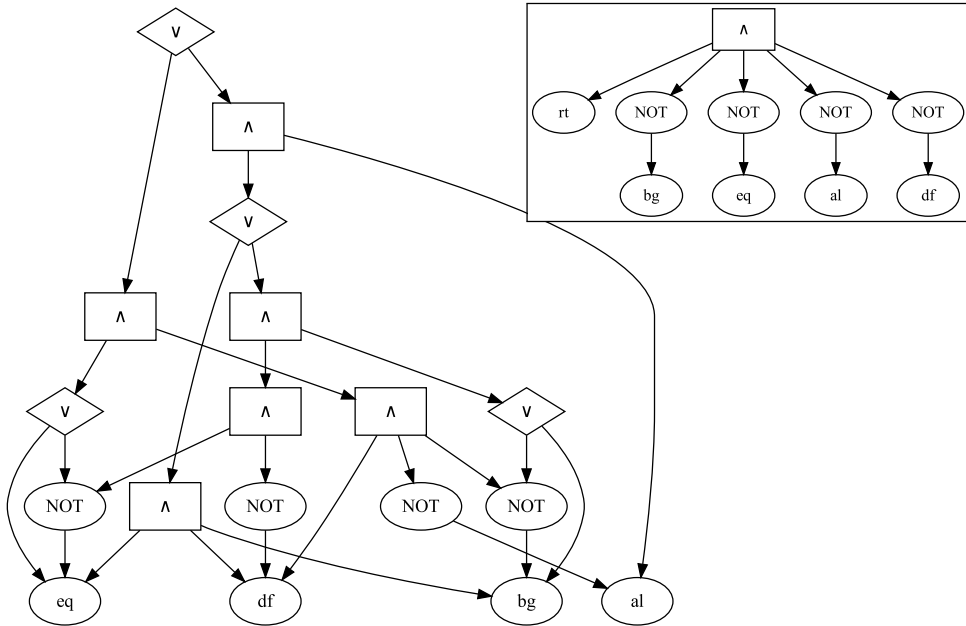


Fig. 7. Compilation step of Example 20 (top right) and Example 22 (left). Models for both theories are compactly encoded as d-DNNF logical circuits. *bg* = burglary, *eq* = earthquake, *al* = alarm, *df* = defective, *rt* = right.

knowledge compiler, because models corresponding to the same total choice may be represented by different nodes and subtrees. Therefore, on this kind of representation we are forced to traverse the circuit in order to retrieve the list of models and the corresponding total choices.

We thus derive from the d-DNNF a circuit where each leaf for a literal *l* is replaced by a set of (partial) models  $\{\{l\}\}$ , disjunctions are replaced by the union of the children and conjunctions correspond to the Cartesian product of the children (Figure 8). Traversing bottom-up such circuit returns the list of models, from which we build a map  $\# : \Omega_{\mathcal{L}} \rightarrow \mathbb{N}$  from total choices to the corresponding number of models, that is,  $\#(\omega) = |\text{MOD}(\mathcal{L}, \omega)|$ . The reciprocal thus defines the normalization constants  $\hat{w}$ . For instance, in Example 22 we obtain  $\#(\omega_1) = \#(\omega_2) = \#(\omega_3) = 1$ ,  $\#(\omega_4) = 2$ , hence  $\hat{w}(\{\text{burglary, alarm, earthquake, right}\}) = \hat{w}(\{\text{burglary, alarm}\}) = \hat{w}(\{\text{earthquake, defective}\}) = 1$ , and  $\hat{w}(\{\text{alarm}\}) = \hat{w}(\{\text{defective}\}) = \frac{1}{2}$ . Note that the total choices that do not appear in the enumeration of models are those that have zero models and thus are inconsistent. Because probabilistic facts are disjoint from the rules' heads, for each model described by the circuit the truth value of the literals corresponding to probabilistic facts defines the corresponding total choice (they can be true iff chosen by  $\omega$ ).

The circuit traversal is a computationally expensive step since it entails enumerating all possible models. Contrary to sums (logical ORs) and products (logical ANDs), the Cartesian product is not a linear-time operation. Therefore, the computational cost of the internal nodes raises with the length of the lists of partial models corresponding to the children. For this reason, the study of novel knowledge compilation techniques that produce a representation where the normalization constants can be obtained from linear time operations is an interesting direction for improvement.

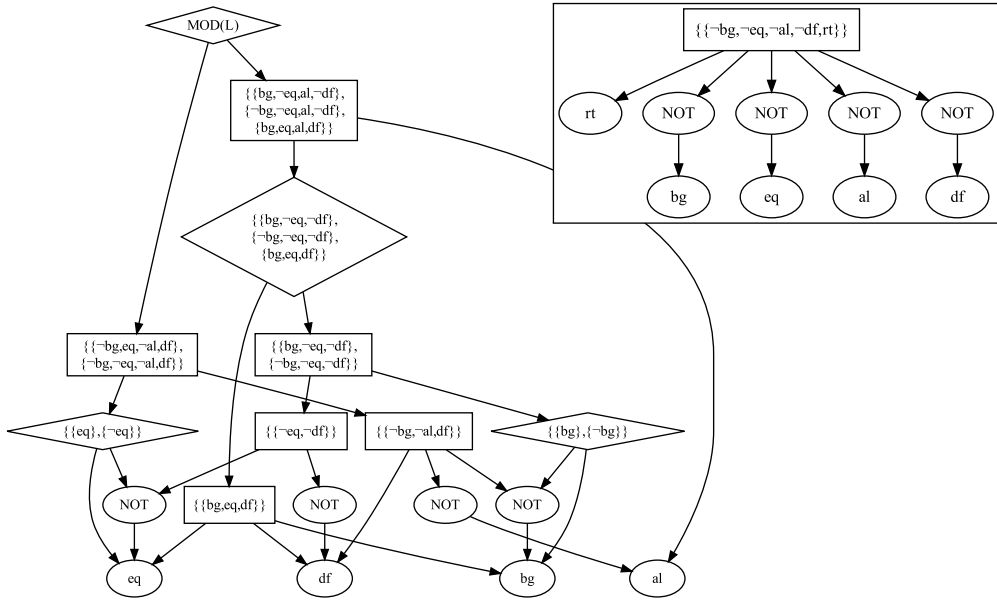


Fig. 8. Enumeration step of Example 20 (top right) and Example 22 models (left). Each node contains the corresponding (partial) models. Diamonds are union nodes, squares are Cartesian products.  $M$  is the complete list of the 5 models of Example 22.

This is a novel problem because SMProbLog semantics applies a different normalization constant per total choice ( $2^n$  with  $n$  facts), while in probabilistic ASP the normalization constant is the same for all models, that is, the weight of all possible worlds. This can be obtained with a single evaluation of the weight of the root of the circuit. Therefore, obtaining the normalization constants from a circuit where node-level operations are computable in linear time would remove this bottleneck, which is introduced because we consider a fundamentally different inference task.

(4) *Evaluation.* As we previously mentioned, to evaluate  $P(q | E = e)$ , with  $\varphi = q \wedge E = e$ , on a program  $\mathcal{L}$  we solve the problem  $\widehat{WMC}_{\mathcal{L}}(\varphi) = \sum_{M \in MOD(\mathcal{L}), M \models \varphi} \hat{w}(M) \prod_{l \in M} w(l)$ . The weights of the leaves in the arithmetic circuit are instantiated according to the query and the given evidence (Section 2). Fierens *et al.* (2011) show that the arithmetic circuit derived from the d-DNNF computes  $WMC_{\mathcal{L}}(\varphi) = \sum_{M \in MOD(\mathcal{L}), M \models \varphi} \prod_{l \in M} w(l)$ . Under a one-to-one correspondence between models and total choices, given  $\omega \in \Omega_{\mathcal{L}}$  and its model  $M_{\omega}$ ,  $WMC(M_{\omega}) = WMC(\omega)$ . Therefore, if the enumeration step is skipped, that is,  $\hat{w}(M_{\omega}) = 1$  for all stable (and well-founded) models  $M_{\omega}$ , then  $WMC(M_{\omega}) = WMC(\omega) = \widehat{WMC}(M_{\omega})$ .

On the contrary, if there are multiple models for some total choices, then  $WMC(M_{\omega}) = WMC(\omega) \neq \widehat{WMC}(M_{\omega})$ , because  $WMC(M_{\omega})$  overcounts the normalized weights of the models corresponding to the same total choice. Figure 9 shows the weight calculation of the root node of the circuit: the root corresponds to all consistent total choices (e.g., Figure 8). Example 22 shows that the original algorithm overcounts the weight when multiple stable models correspond to the same total choice. This because, their weight is not normalized (cfr. Example 36 for the explanation of the 1.25 weight). This raises



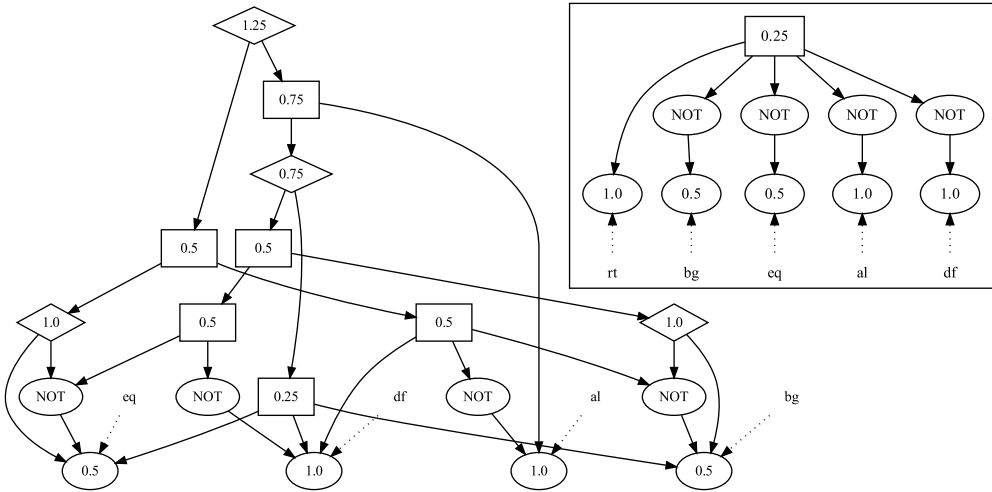


Fig. 9. *WMC* of Example 20 (top right) and Example 22 (left) . Each node contains the corresponding weight, logical variables do not influence the weights. Diamonds are sum nodes, squares are product nodes.

the question as to how the evaluation algorithm has to account for the normalization constants for SMProbLog semantics.

As Figure 8 shows, multiple models can correspond to the same node, therefore it is not possible to associate a normalization constant to a single node. Moreover, the aggregation of the weights in one value per node, as in Figure 9, is no longer possible, because two partial models (e.g.,  $\{bg\}, \{-bg\}$ ) may be conjoined (multiplied) with partial models (weights) that require different normalization constants. Once the weights are aggregated it is not possible to normalize the different components of the weight with different constants. We now describe two evaluation strategies to normalize weights in the evaluation step.

The first is essentially a repetition of the enumeration step: we traverse again the circuit bottom-up and along with each set of atoms we carry the corresponding weight: once a total choice is defined, the weights of each partial model are multiplied and normalized w.r.t. the corresponding number of models obtained from the enumeration step. The drawback of this strategy is that the complexity of the enumeration step is repeated at each query.

The second method is to exploit the efficiency of the *WMC* task on the circuit and compute first the unnormalized weight of the root and then correct it (Algorithm 2). For each total choice  $\omega$  s.t.  $n = \#(\omega) > 1$  and  $\omega \models \varphi$ , the unnormalized evaluation of the circuit overcounts the weight of its models, because they are not multiplied by  $\hat{w}(M_\omega) = \frac{1}{n}$ . Remember that the unnormalized weight of a model is the probability of the corresponding total choice:  $WMC(M_\omega) = WMC(\omega) = P(\omega)$ , therefore the unnormalized weight is such that

$$WMC(M_\omega) = \hat{w}(M_\omega) \cdot WMC(\omega) + (1 - \hat{w}(M_\omega)) \cdot WMC(\omega).$$

Since  $\hat{w}(M_\omega) = \frac{1}{n}$  and  $WMC(\omega) = P(\omega)$ , each weight the of models  $M_\omega \in MOD(\mathcal{L}, \omega)$  is overcounted by

$$(1 - \hat{w}(M_\omega)) \cdot WMC(\omega) = \frac{P(\omega) \cdot (n - 1)}{n}.$$

Removing such weight from the unnormalized one, gives the probability defined by SMProbLog's semantics:

$$\widehat{WMC}(M_\omega) = WMC(M_\omega) - \frac{P(\omega) \cdot (n-1)}{n} = P(\omega) - \frac{P(\omega) \cdot (n-1)}{n} = \frac{P(\omega)}{n} = \hat{P}(M_\omega).$$

The overhead of this method per query lies in the iteration over the total choices with multiple models to retrieve and apply the different normalization constants. On the other hand, there is no additional cost in the operations performed at a node level while traversing the circuit, in contrast to the Cartesian product operations in the first method. Algorithm 2 describes the correction of the weight for any formula  $\varphi$ , where only the models and total choices compatible with  $\varphi$  contribute to  $WMC(\varphi)$  with unnormalized weights. In Section 7 we show experimentally that exploiting the efficiency of the  $WMC$  task results in the evaluation step being significantly faster than the enumeration step, whose complexity reflects on the first method.

Since the circuit describes only the stable models for consistent total choices, in the case of inconsistent programs we derive the probability of an inconsistency by removing from 1 the weight (probability) of the consistent total choices, as described in Section 4.

---

**Algorithm 2** Evaluation step schema

---

**Require:**  $\varphi = q \wedge E = e$

**Output:**  $P(q \mid E = e)$

$w \leftarrow WMC(\varphi)$

**for**  $\omega$  **in**  $\Omega_{\mathcal{L}}$  **s.t.**  $n = \#(\omega) > 1, \omega \models E = e$  **do**

**for**  $M_\omega \in MOD(\mathcal{L}, \omega)$  **s.t.**  $M_\omega \models \varphi$  **do**

$w = w - \frac{P(\omega) \cdot (n-1)}{n}$

**return**  $w$

---

## 6.2 Learning task

In this section we show that the EM learning algorithm implemented in ProbLog2 (Gutmann *et al.* 2011) is correct also for consistent SMProbLog programs. With inconsistent programs, in fact, learning is not possible in a traditional EM setting, because we cannot associate a selection of probabilistic facts to an interpretation of inconsistent atoms. If one of the atoms is inconsistent then all other atoms are inconsistent. Therefore, we can estimate  $\mathbb{P}(\mathcal{L} \models \perp)$  by counting the proportion of inconsistent observations, but we cannot learn how that estimate is distributed across the total choices causing an inconsistency. This information is essential, in this learning setting, to estimate the probability that an atom is chosen. Cussens (2001) extends the EM framework in *stochastic logic programs* (Muggleton *et al.* 1996) to account for failed derivation paths. In future work, investigating the applicability of his *failure adjusted maximization* algorithm to inconsistent probabilistic logic programs might offer a solution to overcome the limits of a traditional EM framework on inconsistent SMProbLog programs. On the other hand, we show that in the case of consistent SMProbLog programs we can learn the probabilities of the program in both total and partial observability.

*Total observability.* The correctness under total observability is guaranteed by the fact that probabilistic facts are independent. In fact, in total observability each interpretation  $I_m \in \mathbb{I}$  observes the truth value of each atom and probabilistic fact of  $\mathcal{L}$ . This case reduces to counting the number of true occurrences of each probabilistic fact in the interpretations  $\mathbb{I}$ . We clarify this by analyzing Equation (1) from Gutmann et al. (2011) which is also implemented in SMProbLog. Let  $\hat{p}_n$  be the estimate for  $p_n :: f_n$ . Let  $\theta_{n,j}^m$  be the  $j$ -th possible grounding substitution for  $f_n$  in the interpretation  $I_m$ . Interpretations  $I_j$  thus represent a possible combination of ground substitutions for the probabilistic facts. Let  $J_n^m$  be the number of such substitutions, and  $Z_n = \sum_{m=1}^M J_n^m$  is the total number of ground instances of  $f_n$  in all training examples:

$$\hat{p}_n = \frac{1}{Z_n} \sum_{m=1}^M \sum_{j=1}^{J_n^m} \delta_{n,j}^m \quad \text{where} \quad \delta_{n,j}^m = \begin{cases} 1 & \text{if } f_n \theta_{n,j}^m \in I_m; \\ 0 & \text{else.} \end{cases} \quad (1)$$

The outermost sum ranges over all observed interpretations and the innermost sum counts each different observation of a ground instance of  $f_n$ . Each number of ground substitutions observed is then divided by the total number of ground substitutions. The estimate  $\hat{p}_n$  is thus defined as the proportion of ground substitutions observed over all interpretations observed. This is correct for SMProbLog programs because the observation of a probabilistic fact remains independent of the rest of the program for each interpretation.

*Partial observability.* The correctness under partial observability is guaranteed by the fact that the parameter updates rely on probabilistic inference, which is performed under SMProbLog's semantics and thus returns correct probability estimates. In the partially observable case, where each interpretation  $I_m$  observes the truth value of a subset of  $\mathcal{L}$ , the parameter update iteration relies on probabilistic inference to update the likelihood of a fact given  $I_m$ . At each iteration  $k$  the parameters from the previous iteration  $k-1$  are used in  $\mathcal{L}$  to compute the conditional expectation of the parameter given the interpretations, until convergence. The number of observations  $\delta_{n,j}^m$  in ((1)) is replaced by the conditional expectation under the current model  $\mathbb{E}[\delta_{n,j}^m | I_m]$ :

$$\hat{p}_n = \frac{1}{Z_n} \sum_{m=1}^M \sum_{j=1}^{J_n^m} \mathbb{E}[\delta_{n,j}^m | I_m]. \quad (2)$$

To compute the conditional expectation values we query the model for  $\mathbb{P}(f_n)$  at each step under the current parameters estimate. This means that the normalization w.r.t. multiple stable models is incorporated in the conditional probabilities computed by the inference task, therefore also the parameter estimate is normalized w.r.t. the different stable models that may correspond to a given partial observation. Since the algorithmic procedure for learning in SMProbLog remains the same as ProbLog2, except from the inference step, Example 6 is a valid example also for learning from a partial interpretation in SMProbLog.

## 7 Experiments

The goal of our experiments is to establish the feasibility of our approach, and to answer the following questions:

- (Q1) What is the computational cost of the different steps in the pipeline for SMProbLog semantics on the inference task?
- (Q2) Given a set of observations of accepted arguments generated from a known ProbLog program, can we learn the degrees of belief of the agent that modeled the original program?
- (Q3) How does the inference computational cost reflect on learning?

*Q1. Inference.* We consider a variation of a typical PLP example where a set of people has a certain probability of having asthma or being stressed, and stress leads with some probability to smoking:

$$\begin{aligned} 0.1 &:: \text{asthma}(X) \leftarrow \text{person}(X). \\ 0.3 &:: \text{stress}(X) \leftarrow \text{person}(X). \\ 0.4 &:: \text{smokes}(X) \leftarrow \text{stress}(X). \end{aligned}$$

People are related by an influence relation: if a person smokes and influences to some extent another one, then the other person will smoke, and if someone smokes there is a probability to suffer from asthma. If someone suffers from asthma then the person does not smoke (an example of a cycle with negation).

$$\begin{aligned} &\text{smokes}(X) \leftarrow \text{influences}(Y, X), \text{smokes}(Y). \\ 0.4 &:: \text{asthma}(X) \leftarrow \text{smokes}(X). \\ &\neg \text{smokes}(X) \leftarrow \text{asthma}(X). \end{aligned}$$

With this last rule we add a cycle through negation to a classical standard benchmark used in the statistical relational learning community. This benchmark is often used to study the scaling capabilities of frameworks by growing the size of the domain of people. Let  $R$  be the aforementioned rules, we consider examples with an increasing number of people and relationships:

- $t_1 = R \cup \{\text{person}(1). \text{person}(2). 0.3 :: \text{influences}(1, 2). 0.6 :: \text{influences}(2, 1).\}$
- $t_2 = t_1 \cup \{\text{person}(3).\}$
- $t_3 = t_2 \cup \{\text{person}(4).\}$
- $t_4 = t_3 \cup \{0.2 :: \text{influences}(2, 3).\}$
- $t_5 = t_4 \cup \{0.7 :: \text{influences}(3, 4).\}$
- $t_6 = t_5 \cup \{0.9 :: \text{influences}(4, 1).\}$

Figure 10 shows the running time of SMProbLog on the different benchmarks for the queries  $\text{query}(\text{smokes}(X))$ . and  $\text{query}(\text{asthma}(X))$ . The cost of grounding (step 1) is negligible, hence we omit it from the plot. This experiment answers question Q1: as expected, the enumeration step dominates the running time of SMProbLog and it is directly related to the size of the circuit. A larger circuit is likely to present a higher number of Cartesian products in the enumeration step, hence a higher computational cost is expected. This experiment confirms empirically the expected exponential cost of computing weights for the SMProbLog semantics on a standard d-DNNF, in contrast with the polynomial time required for the traditional *WMC* task. At the same time, it provides an example of a problem where the optimization of the evaluation step presented in Section 6.1 prevents the repetition of the cost of the enumeration step. The cost of the evaluation step is in fact consistently lower than the enumeration step. Finally, we remark

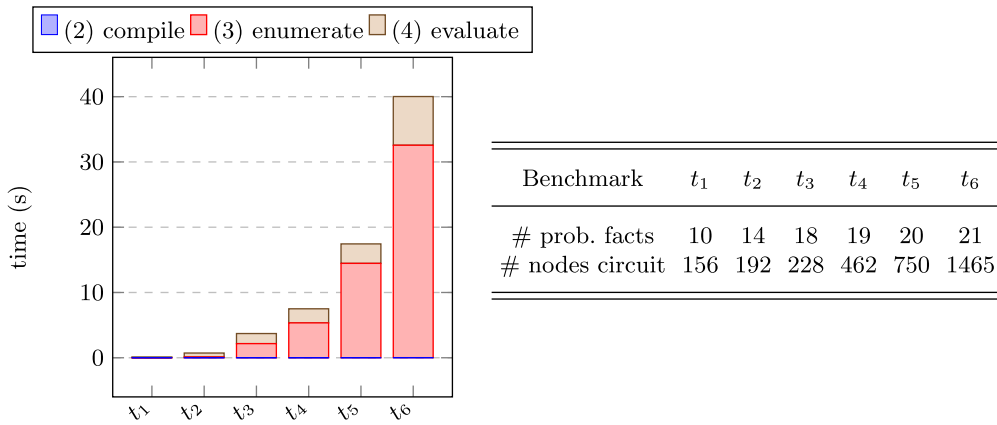


Fig. 10. Inference time on benchmarks.

that the approach based on knowledge compilation allows us to answer an increasing number of (ground) queries: 4 for  $t_1$ , 6 for  $t_2$ , and 8 for the remaining benchmarks. This affects the evaluation time, but the grounding, compilation, and enumeration steps are executed once, regardless of the number of queries.

*Q2. Learning.* We answer question 2 by considering a dataset of 283 argument graphs (Stede et al. 2016) and by deriving from each annotation of the dataset a bipolar argument graph. We attach to arguments and relations random probabilities to reflect an agent's belief in each, as we did for the running example, Example 8. The graphs contain from 4 to 11 nodes (average 5.3) with an average of 1.4 attacks and 3 supports per graph.

Given this dataset of probabilistic argument graphs, we test learning as follows. For each argumentation graph modeled in a program  $\mathcal{L}$  we obtain the model to be learned  $\mathcal{L}^*$  by replacing each probability with a (randomly initialized) learnable parameter  $t(\cdot)$ , similarly to Example 6. The programs  $\mathcal{L}^*$  contain on average 10 learnable parameters divided between probabilistic facts (arguments' biases) and probabilistic rules (relations' belief). We use  $\mathcal{L}$  to sample  $n$  observations of the argumentation graph by means of the sampling tool from ProbLog2. Then, we use  $\mathcal{L}^*$  to learn the original probabilities from the samples. We evaluate the accuracy of the learned parameter by considering the mean absolute error (MAE) of the learned probabilities in  $\mathcal{L}^*$  compared to the original ones in  $\mathcal{L}$ . We run the learning test with increasing size of the samples, that is,  $n \in \{50, 100, 150, 200, 250, 300\}$ , and with an upper bound of 100 EM iterations.

Note that we are in the case of partial observability, as the parameters are attached to the predicates *bias*, *arg<sub>pos</sub>*, and *arg<sub>neg</sub>*, but we observe only the outcome in the form of the predicate *arg*. We evaluate the quality of the learned programs with the MAE, summarized in Figure 11. From the experiments we observe that the MAE decreases with the increase of the sample size, that is, the more observations from the original distribution are provided, the higher is the accuracy in learning the original probabilities. At the same time when at least 100 samples are provided, the MAE drops below 10%. Therefore, we can answer positively to our question about learnability of beliefs in argument graphs with SMProbLog.




Fig. 11. Mean absolute error by number of samples.




Fig. 12. Mean running time by number of parameters and size of circuit (dashed line) on increasing number of observations (colored lines).

*Q3. Learning vs. inference.* Under the same experimental setup of *Q2*, we now analyze the time required to learn the parameters with respect to the cost of inference, which is part of the EM learning algorithm (Equation (2)). There are two dimensions that influence the time required for learning with respect to inference, which we represent in Figure 12. The first is the cost of updating the probabilities for an increasing number of observations. In Figure 12 we show the running time (left scale) of EM learning on increasing sample sizes in 100, 150, 200, 250, 300. The second is the intrinsic cost of inference with respect to the complexity of the program. We measure the complexity of the program by measuring the size of the compiled circuit, represented in Figure 12 with a dashed line (right scale).

The experiment shows that the second dimension is the main factor in determining the time required for learning, which follows the same trend of inference shown in Figure 10. Since we replace all probabilistic facts with learnable parameters, the x-axis of the two figures are comparable. The dashed line represents the average size of the circuit for a set of programs with a given number of learnable parameters (x-axis). As expected from Equation (2), the experiment shows a strong correlation between the inference cost on an increasingly large circuit and the learning time. In particular, we observe that the average learning time with 17 parameters decreases from 15, because in the dataset the average size of the circuits with 17 parameters is also smaller. The first dimension of the experiment shows that on smaller models the difference between the sample sizes (colored lines) is negligible, while on more complex models the differences spread out. This is plausible because on more complex models the cost of inference is higher and thus the repetition of a larger cost over more examples becomes more evident. However, the impact of a higher number of observation on the total time required for learning is clearly inferior to a bigger (more complex) logical circuit.

## 8 Related work

In this section we discuss the related work with respect to PLP and probabilistic argumentation.

### 8.1 Probabilistic logic programming.

Both three-valued well-founded model semantics and stable model semantics have been considered to extend traditional PLP framework's semantics: Hadjichristodoulou and Warren (2012) present an extension of PRISM based on three-valued well-founded models. When considering general normal logic programs, well-founded semantics (Gelder et al. 1991) is a three-valued semantics, that is, logical atoms can be true, false, or undefined. A three-valued interpretation is a pair  $I = (T, F)$  where  $T$  and  $F$  are disjoint subsets of the Herbrand base  $HB(\mathcal{L})$  of the program  $\mathcal{L}$ . The (ground) atoms in  $T$  (resp.  $F$ ) are interpreted as *true* (resp. *false*) in  $I$ . The remaining atoms in  $HB(\mathcal{L}) \setminus (T \cup F)$  are said to be *undefined*. A three-valued model is thus a three-valued interpretation which satisfies all the rules of the program. Our approach is different because with stable model semantics we assign a truth value to the non-deterministic choices that well-founded semantics would label as undefined. At the same time we use the third value to denote the absence of stable models.

On the other hand, the application of stable model semantics in PLP is represented by probabilistic answer set programming (ASP) (Cozman and Mauá 2017; 2020). Table 3 summarizes the differences between the frameworks, which emerge on the definition of the probability distribution and the definition of model of a possible world.

*Relation with probabilistic ASP.* The key difference between SMProbLog and probabilistic ASP frameworks is that our semantics is based on the distribution semantics, that is, we define a probability distribution over models from a probability distribution over total choices. On the contrary, probabilistic ASP languages such as P-Log (Baral et al. 2009) and LP<sup>MLN</sup> Lee and Wang (2016) use the probability labels to directly define a

Table 3. *PLP frameworks comparison. n.a. = not applicable. 2\* = the logic semantics is two-valued stable models, but we introduce a third value for total choices with 0 stable models.*

Framework	Distribution Model semantics	Logic values	Models per total choice
ProbLog2 (Fierens <i>et al.</i> 2011)	✓	Well-founded	2
WF-PRISM (Hadjichristodoulou and Warren 2012)	✓	Well-founded	3
P-Log (Baral <i>et al.</i> 2009)	✗	Stable models	2
LP <sup>MLN</sup> (Lee and Wang 2016)	✗	Stable models	2
SMProbLog	✓	Stable models	2*
			≥ 1

(globally normalized) probability distribution over the models of a (derived) program. This difference allows us to preserve the marginal probability of independent facts in the joint distribution in the presence of multiple stable models, and to measure the probability of an inconsistency in a program. In fact, following the probabilistic ASP approach can give counterintuitive results as we show in Example 36.

#### Example 36

In Example 22, if we associate to each stable model  $M_\omega$  the probability of the corresponding total choice  $P(\omega)$ , the sum of the probabilities of the stable models is  $5 \cdot 0.25 = 1.25$ . A normalization w.r.t. all answer sets, as P-Log does, leads to a probability distribution over stable models  $\hat{P}'(M_\omega) = 0.25/1.25 = 0.2$  for all  $M_\omega$ . Then the probability of a query *burglary* is  $\mathbb{P}(\textit{burglary}) = \hat{P}'(M_{\omega_1}) + \hat{P}'(M_{\omega_2}) = 0.4$ . This means that the marginal probability of *burglary* in the joint distribution is lower than the prior 0.5 despite no epistemic influence is defined on *burglary*. In our approach, the marginal probability of that fact being true under the final distribution equals the fact's label, and is thus directly interpretable. Similarly, when some possible worlds are inconsistent, the corresponding probability mass is not extended to a model, but rather a global normalization redistributes uniformly such mass to the valid models. For this reason, in Example 20 a global normalization would assign all the probability mass to  $\omega_4$ , that is,  $\hat{P}'(\omega_4) = \frac{0.25}{0.25} = 1$ . This because  $\{\textit{right}\}$  is the only possible stable model hence its weight, 0.25, is also the sum of all possible stable models. Therefore, all atoms but *right* would be interpreted false and *right* would be interpreted true with probability 1, thus making the presence of logical inconsistencies unclear.

#### Example 37

In Example 17 we have a similar situation: by weighing each stable model with the probability labels of the literals in the model we have that the sum of the weights of the models is 1.42 and that  $\hat{P}'(M_{\omega_1}) = \frac{0.12}{1.42} = 0.085$ ,  $\hat{P}'(M_{\omega_2}) = \frac{0.18}{1.42} = 0.1268$ ,  $\hat{P}'(M_{\omega_3}) = \frac{0.28}{1.42} = 0.1972$ ,  $\hat{P}'(M_{\omega_4}^1) = \hat{P}'(M_{\omega_4}^2) = \frac{0.42}{1.42} = 0.2958$ . Therefore, the probabilities of *too\_expensive\_X* and *too\_noisy\_Y* are respectively  $\hat{P}'(M_{\omega_2}) + \hat{P}'(M_{\omega_4}^1) + \hat{P}'(M_{\omega_4}^2) = 0.7183$  and  $\hat{P}'(M_{\omega_3}) + \hat{P}'(M_{\omega_4}^1) + \hat{P}'(M_{\omega_4}^2) = 0.7887$ . Therefore, the probability of probabilistic facts in the inferred distribution increases w.r.t. the original declared value. This happens with a globally normalized definition of the probability of the models and multiple stable models per total choice.






Fig. 13. Possible colorings of Example 38.

In this paper we focus on computing exact probabilities as Baral *et al.* (2009) and Lee and Wang (2016) rather than defining an interval of probabilities for each atom as in credal semantics (Lukasiewicz 1998). Cozman and Mauá (2020) remarks the difference between a credal semantics approach and a distribution semantics approach like ours. They emphasize that in distribution semantics a probabilistic fact  $p :: f$ . the fact  $f$  is imposed with probability  $p$  and discarded with probability  $1 - p$ , while their approach distinguishes between taking  $f$  with probability  $p$  and  $\neg f$  with probability  $1 - p$ . Moreover, credal semantics are defined only for consistent logic programs, as in Azzolini *et al.* (2022). Recent work (Rocha and Cozman 2022) extends credal semantics to L-Credal semantics to handle inconsistencies. We illustrate the differences by means of the graph coloring example from Cozman and Mauá (2016), other comparisons are also present in Rocha and Cozman (2022).

#### Example 38

Consider the following program:

---

```

coloredBy(V, red) ← ∼coloredBy(V, yellow), ∼coloredBy(V, green), vertex(V).
coloredBy(V, yellow) ← ∼coloredBy(V, red), ∼coloredBy(V, green), vertex(V).
coloredBy(V, green) ← ∼coloredBy(V, red), ∼coloredBy(V, yellow), vertex(V).
noClash ← ∼noClash, edge(V, U), coloredBy(V, C), coloredBy(U, C).
edge(1, 4).   edge(2, 1).   edge(2, 4).   edge(3, 5).   edge(4, 3).
edge(1, 3).  0.5 :: edge(4, 5).  coloredBy(2, red).  coloredBy(5, green).

```

---

Credal semantics here define probability intervals based on the stable models of the two cases where  $edge(4,5)$  is present (one model) or not (two models). The possible colorings are represented in Figure 13. Credal semantics define a lower ( $\mathbb{P}$ ) and upper ( $\overline{\mathbb{P}}$ ) probability for each node:  $\mathbb{P}(coloredBy(1, yellow)) = 0$ ,  $\overline{\mathbb{P}}(coloredBy(1, yellow)) = \frac{1}{2}$ ,  $\mathbb{P}(coloredBy(4, yellow)) = \frac{1}{2}$ ,  $\overline{\mathbb{P}}(coloredBy(4, yellow)) = 1$ , and  $\mathbb{P}(coloredBy(3, red)) = \overline{\mathbb{P}}(coloredBy(3, red)) = 1$ . On the contrary with our semantics we compute exact values:  $\mathbb{P}(coloredBy(1, yellow)) = 0.25$ ,  $\mathbb{P}(coloredBy(4, yellow)) = 0.75$ ,  $\mathbb{P}(coloredBy(3, red)) = 1$ . These values derive from the probability of 0.5 of the coloring with  $edge(4,5)$  and the probability of  $\frac{0.5}{2}$  of each of the two possible coloring without  $edge(4,5)$ .

*Relations with traditional PLP.* Our approach generalizes traditional PLP frameworks on programs without function symbols: when a program defines total choices corresponding to exactly one two-valued well-founded model, then the semantics agree on both models

and probability. As for the model, if a normal logic program has a total two-valued well-founded model, then the model is the unique stable model (Gelder *et al.* 1991). Thus, in this case  $|M(\omega)| = 1$  for each total choice  $\omega$ , from which it follows that  $\hat{P}(M_\omega) = P(\omega)$  for the single stable model  $M_\omega \in MOD(\mathcal{L}, \omega)$  ( $MOD(\mathcal{L}, \omega) = \{M_\omega\}$ ). This means that the probability of the model is the probability of the corresponding total choice as in (probabilistic) two-valued well-founded semantics. The one-to-one correspondence between total choices and stable models also guarantees that  $\mathbb{P}((\emptyset, \emptyset)) = 0$ .

For this reason, our pipeline is also a generalization of ProbLog2's pipeline. ProbLog2 (Fierens *et al.* 2011) reduces the probabilistic inference task to a *WMC problem* (Cadoli and Donini 1997) in three steps:

1. Compute the relevant grounding, that is, ground only the part of the program necessary to answer the query.
2. Convert the ground rules into an equivalent boolean formula (CNF).
3. Compile the boolean formula into an arithmetic circuit to efficiently compute the weighted model count of the formula.

The grounding procedure in ProbLog2 computes the *relevant* ground program w.r.t. the given queries  $Q$  and evidence  $E$ . The relevant ground program is obtained by applying SLD resolution to prove all atoms in  $Q \cup E$ . However, in SMProbLog an atom can be derived from a logical choice whose atoms are not directly justified by some (probabilistic) fact, therefore in this case meaningful parts of the program would not be included in the relevant ground program. This is why we use a standard bottom-up technique. Aziz *et al.* (2015) apply stable model counting techniques in ProbLog (ASPProbLog). In particular, they replace the DSHARP compiler (Muise *et al.* 2012) used to implement step 3 with a stable model knowledge compiler (an extension of DSHARP itself). Step 2 is also modified, since the CNF conversion is skipped and a simple representation of the ground rules is provided to the stable model knowledge compiler. However, ASPProbLog does not change the semantics of ProbLog, and invalid programs according to the original semantics are still rejected. Moreover, ASPProbLog does not implement the syntactical features of ProbLog2 on which our approach to argumentation is based, such as negation in the heads or annotated rules.

Distribution semantics has been proven by Sato (1995) to be well-defined for definite programs with function symbols, and later Riguzzi (2015) proves this for the case of normal programs with two-valued well-founded models. Studying the semantics of SMProbLog programs with function symbols is thus an interesting direction for future work.

## 8.2 Probabilistic argumentation

Our approach is novel as previous work considers different semantics and reasoning techniques for probabilistic argument graphs. Table 4 summarizes these differences: the choice between interpreting probabilities as in the constellation approach or as in the epistemic approach, approaching the problem from PLP modeling a single joint probability distribution as in Bayesian networks, whether it can reason about marginal and conditional queries about the probability of arguments, and the implementation in a framework where the probabilities of the graph can be learned. Moreover, our framework is the

Table 4. *Argumentation frameworks comparison.*  $\sim$  denotes partial support

Framework	Epistemic	PLP	Conditionals and marginals	Implementation
QBAF (Baroni et al. 2019)	✗	✗	✗	✗
Prob. argument graphs (Li et al. 2011)	✗	✗	✗	✓
MetaProlog (Mantadelis and Bistarelli 2020)	✗	✓	✗	✓
Epistemic graphs (Hunter et al. 2020)	✓	✗	✗	✓
L-Credal (Rocha and Cozman 2022)	✓	✓	~	✗
SMProbLog	✓	✓	✓	✓

only one providing an implementation that can learn from observations (of accepted arguments).

*Epistemic.* We follow the epistemic interpretation of probabilities (Hunter et al. 2020; Hunter 2013) as a direct measure of the belief of an agent in arguments. Section 5 already discussed the main differences with previous work. We use fine-grained (gradual) evaluations of relations and base scores for arguments in a bipolar setting similarly to quantified bipolar argument frameworks (QBAFs) (Baroni et al. 2019). Similarly to our framework, in a QBAF a set of arguments contains both an attack and support relation, and defines a base score  $\tau$ , which corresponds in our case to the prior probability (belief)  $P_A$ , as Baroni et al. (2019) describes it: “. . . the nature and meaning of the base score [. . .] corresponds to an assessment of arguments which precedes the consideration of the relations of attack and support with other arguments.” In our case the joint probability distribution is the result of this consideration. QBAF is a restricted setting compared to our framework since we also define a score (following the terminology of Baroni et al. (2019)) for the relations  $R^+$  and  $R^-$ , besides considering set-attacks. Most importantly, we do this with probabilistic semantics.

*PLP.* At the same time, we propose a mapping from probabilistic argument graphs to probabilistic logic programs, which provides a novel semantics for probabilistic argument graphs. We consider a single probability distribution in a Bayesian style rather than families of distributions. Previous work in fact focused on reasoning about the properties of families of probability distributions that are consistent with the argument graph and additional constraints (*epistemic graphs*) (Hunter et al. 2020).

Despite following the epistemic approach, our framework has similarities with the constellations approach. The definition of the distribution over subgraphs in the constellations approach is similar to the distribution over subprograms  $F \cup R$  of distribution semantics. In fact, the distribution modeled by the probabilistic logic program is determined by the probability that an argument is included in a possible world, similar to the inclusion or exclusion of nodes and edges in the subgraphs considered in the constellations approach. However, in order to determine whether an argument is true or not, we rely on the stable model semantics for logic programs. In the constellations approach, on the other hand, the admissible arguments are defined by the classical extension-based

semantics (Dung 1995). Stable models are equivalent to stable extensions for basic arguments graphs (Dung 1995). We consider however more sophisticated relations between arguments that cannot be encoded in a traditional (probabilistic) abstract argumentation framework, for example, support (Amgoud *et al.* 2004) or set-attacks (Nielsen and Parsons 2006). For this reason, our approach does not fit any of the previous categorizations about the connections between logic (programming) and argumentation (Besnard *et al.* 2020).

Modeling (deterministic) argumentation problems by means of logic programming dates back to the foundational work of Dung (1995). Dung (1995) proposes a method to define meta-interpreters for argument systems encoding the extension-based semantics as logic rules. A change in argumentation semantics thus requires to encode in the meta-interpreter a new reasoning technique. This is the approach taken in Mantadelis and Bistarelli (2020), while we directly compute the stable models semantics for logic programs on a mapping from argument graphs to probabilistic logic programs.

Wu *et al.* (2009) propose a mapping for Dung's abstract argumentation frameworks to logic programs, therefore they do not address the probabilistic settings as well as other extensions to the original framework considered in this paper. Wu *et al.* (2009) focus on complete labelings and propose a translation where given an argument graph  $(A, R)$ , given a set of attacks of the form  $(b_i, a) \in (A, R)$ ,  $i \in \{1, \dots, n\}$  a logic rule of the form  $a \leftarrow \sim b_1, \dots, \sim b_n$ . is added to the logic program translation. This is a similar principle followed in our encoding, where attacks  $(b_i, a)$  with  $P_R((b_i, a)) = p_i$  are translated to rules of the form  $p_i :: \neg a \leftarrow b_i$ . In fact, the rewriting of negation in the head produces rules  $a \leftarrow a_{pos}, \sim a_{not}$  plus  $p_i :: a_{not} \leftarrow b_i$ . Therefore, in both cases all attackers  $b_i$  must be false in order to consider  $a$  acceptable. However, there is a difference between negation in the head (and the corresponding translation) and negation in the body. Attacks directly encoded with negation in the body do not exclude that the two choices can be true at the same time because of external justification, but exclude that both are false. Vice versa, an encoding with negation in the head excludes that both choices are true at the same time, but not that both can be false (e.g., in Example 23).

### Example 39

Consider the difference between a program with a loop with negation in the body (A) and one with negation in the head (B).

A	B
$0.6 :: \text{too\_expensive\_X}.$	$0.6 :: \text{too\_expensive\_X}.$
$0.7 :: \text{too\_noisy\_Y}.$	$0.7 :: \text{too\_noisy\_Y}.$
$\text{stay\_at\_Y} \leftarrow \text{too\_expensive\_X}.$	$\text{stay\_at\_Y} \leftarrow \text{too\_expensive\_X}.$
$\text{stay\_at\_X} \leftarrow \text{too\_noisy\_Y}.$	$\text{stay\_at\_X} \leftarrow \text{too\_noisy\_Y}.$
$\neg \text{stay\_at\_X} \leftarrow \text{stay\_at\_X}.$	$\text{stay\_at\_X} \leftarrow \sim \text{stay\_at\_Y}.$
$\neg \text{stay\_at\_Y} \leftarrow \text{stay\_at\_X}.$	$\text{stay\_at\_Y} \leftarrow \sim \text{stay\_at\_X}.$

Program A is the program of Example 23. Program B is a similar encoding with negation in the body. In this case the possible worlds  $\omega_2$  and  $\omega_3$  have the same models of A. However, for  $\omega_1 = \{\}$  where program A has an empty model, program B has the models

Table 5. Possible models comparison of Example 39

Total choice	Model A	Model B
$\omega_1 = \{\}$	$\{\}$	$\{\{stay\_at\_Y\}, \{stay\_at\_X\}\}$
$\omega_2 = \{too\_expensive\_X\}$	$\{\omega_2 \cup \{stay\_at\_Y\}\}$	$\{\omega_2 \cup \{stay\_at\_Y\}\}$
$\omega_3 = \{too\_noisy\_Y\}$	$\{\omega_3 \cup \{stay\_at\_X\}\}$	$\{\omega_3 \cup \{stay\_at\_X\}\}$
$\omega_4 = \{too\_expensive\_X, too\_noisy\_Y\}$	$\{\omega_4 \cup \{stay\_at\_Y\}, \omega_4 \cup \{stay\_at\_X\}\}$	$\{\omega_4 \cup \{stay\_at\_Y, stay\_at\_X\}\}$

$\{\{stay\_at\_Y\}, \{stay\_at\_X\}\}$ . While program B forces a choice when the counterarguments are not accepted, program A allows no choice. At the same time, program B allows on  $\omega_4$  to infer a stay at both hotels from the counterarguments for each, that is, the model  $\{too\_expensive\_X, too\_noisy\_Y, stay\_at\_X, stay\_at\_Y\}$ . Table 5 compares the models for each total choice.

Past work on the connections with logic programming beyond basic abstract argumentation frameworks are limited to the deterministic setting, except for the work by Rocha and Cozman (2022). They adopt the translation by Wu et al. (2009) to encode assumption based argumentation frameworks (Bondarenko et al. 1993) in a probabilistic logic program. Caminada et al. (2015) explore equivalence relations for more expressive deterministic argumentation frameworks, by framing logic programming semantics in terms of argumentation semantics. Alfano et al. (2020) consider the opposite approach and follow in a *deterministic* setting a similar approach to ours for *probabilistic* argumentation. In fact, a simple but general logical framework is shown to be able to capture, in a systematic and succinct way, the different features of several argumentation frameworks under different argumentation semantics. The authors remark how the flexibility of a logic programming approach encourages the study of more extensions and can be used for better understanding the semantics of extended argumentation frameworks.

*Marginals and conditionals.* With our methodology it is possible to apply general PLP reasoning techniques to probabilistic argument graphs, such as marginal probability computation (Example 27), conditioning over evidence (Example 28), and parameter learning (Section 6.2). This results in a modular, expressive, extensible framework reflecting the dynamic nature of beliefs in an argumentation process. On the contrary, argumentation systems like epistemic graphs require (ad-hoc) reasoning algorithms that do not follow the traditional probabilistic reasoning techniques. In the framework by Rocha and Cozman (2022) conditioning on atoms that appear negated in a body is deferred to future work.

*Implementation.* In Section 6 we presented a system that computes SMProbLog semantics, while in the past the focus has been on developing new argumentation semantics rather than tools to compute them. Moreover, we provide an implementation of the EM learning algorithm for probabilistic parameters whose application to probabilistic argumentation problems is novel.

## 9 Conclusion

We proposed a novel approach to epistemic probabilistic argumentation in two points. First, the interpretation of a probabilistic argument graph as a probabilistic graphical model defining a joint probability distribution over arguments. Second, we presented a novel syntactical translation of argument attacks to logic programming rules based on negation in the head and the inhibition effect.

Our method regards probabilistic argument graphs as the description of how the prior beliefs (biases) interact with each other to define marginal and conditional beliefs that are coherent with the argumentative (logic) structure. This allows us to exploit the structure of conditional (in)dependencies of arguments and apply the traditional PLP inference methods to the argumentative representation of subjective beliefs, such as learning, conditional queries, and MPEs.

Approaching probabilistic argumentation from a PLP perspective stresses the limiting assumptions of PLP frameworks when (probabilistic) normal logic programs are concerned. For this reason in this paper we proposed a new PLP system, SMProbLog, based on a combination of the classical distribution semantics for probabilistic logic programs with stable model semantics.

SMProbLog generalizes previous work on distribution semantics for programs without function symbols, supporting inference and parameter learning tasks for a wider class of (probabilistic) logic programs. The experiments in this regard show that the existing PLP inference techniques can be used, but under the new semantics not all the necessary information can be computed efficiently. Novel knowledge compilation techniques allowing polynomial-time querying under SMProbLog's semantics would provide a significant improvement for both the inference and learning tasks, therefore future research on this topic represents an interesting and useful development of this work.

## Competing interests

The authors declare none.

## References

- ALFANO, G., GRECO, S., PARISI, F. AND TRUBITSYNA, I. 2020. On the semantics of abstract argumentation frameworks: A logic programming approach. *Theory and Practice of Logic Programming* 20, 5, 703–718.
- AMGOUD, L., BEN-NAIM, J., DODER, D. AND VESIC, S. 2017. Acceptability semantics for weighted argumentation frameworks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19–25, 2017*, C. Sierra, Ed. ijcai.org, 56–62.
- AMGOUD, L., CAYROL, C. AND LAGASQUIE-SCHIEX, M. 2004. On the bipolarity in argumentation frameworks. In *10th International Workshop on Non-Monotonic Reasoning (NMR 2004), Whistler, Canada, June 6–8, 2004, Proceedings*, J. P. Delgrande and T. Schaub, Eds. 1–9.
- AZIZ, R. A., CHU, G., MUISE, C. J. AND STUCKEY, P. J. 2015. Stable model counting and its application in probabilistic logic programming. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA*, B. Bonet and S. Koenig, Eds. AAAI Press, 3468–3474.

- AZZOLINI, D., BELLODI, E. AND RIGUZZI, F. 2022. Statistical statements in probabilistic logic programming. In *Logic Programming and Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022, Genova, Italy, September 5-9, 2022, Proceedings*, G. Gottlob, D. Inglezian and M. Maratea, Eds. Lecture Notes in Computer Science, vol. 13416. Springer, 43–55.
- BARAL, C., GELFOND, M. AND RUSHTON, J. N. 2009. Probabilistic reasoning with answer sets. *Theory and Practice of Logic Programming* 9, 1, 57–144.
- BARONI, P., RAGO, A. AND TONI, F. 2019. From fine-grained properties to broad principles for gradual argumentation: A principled spectrum. *International Journal of Approximate Reasoning* 105, 252–286.
- BERRY, D. A. 2006. Bayesian clinical trials. *Nature Reviews Drug Discovery* 5, 1, 27–36.
- BESNARD, P., CAYROL, C. AND LAGASQUIE-SCHIEX, M. 2020. Logical theories and abstract argumentation: A survey of existing works. *Argument Comput.* 11, 1-2, 41–102.
- BONDARENKO, A., TONI, F. AND KOWALSKI, R. A. 1993. An assumption-based framework for non-monotonic reasoning. In *Logic Programming and Non-monotonic Reasoning, Proceedings of the Second International Workshop, Lisbon, Portugal, June 1993*, L. M. Pereira and A. Nerode, Eds. MIT Press, 171–189.
- CADOLI, M. AND DONINI, F. M. 1997. A survey on knowledge compilation. *AI Communications* 10, 3-4, 137–150.
- CAMINADA, M., SÁ, S., ALCÁNTARA, J. F. L. AND DVORÁK, W. 2015. On the equivalence between logic programming semantics and argumentation semantics. *International Journal of Approximate Reasoning* 58, 87–111.
- CAYROL, C. AND LAGASQUIE-SCHIEX, M. 2005. Graduality in argumentation. *Journal of Artificial Intelligence Research* 23, 245–297.
- CHARNIAK, E. 1991. Bayesian networks without tears. *AI Magazine*, 12, 4, 50–63.
- CHOI, A., KISA, D. AND DARWICHE, A. 2013. Compiling probabilistic graphical models using sentential decision diagrams. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 12th European Conference, ECSQARU 2013, Utrecht, The Netherlands, July 8-10, 2013. Proceedings*, L. C. van der Gaag, Ed. Lecture Notes in Computer Science, vol. 7958. Springer, 121–132.
- COZMAN, F. G. AND MAUÁ, D. D. 2016. The structure and complexity of credal semantics. In *Proceedings of the 3rd International Workshop on Probabilistic Logic Programming co-located with 26th International Conference on Inductive Logic Programming (ILP 2016), London, UK, September 3, 2016*, A. Hommersom and S. A. Abdallah, Eds. CEUR Workshop Proceedings, vol. 1661. CEUR-WS.org, 3–14.
- COZMAN, F. G. AND MAUÁ, D. D. 2017. On the semantics and complexity of probabilistic logic programs. *Journal of Artificial Intelligence Research*, 60, 221–262.
- COZMAN, F. G. AND MAUÁ, D. D. 2020. The joy of probabilistic answer set programming: Semantics, complexity, expressivity, inference. *International Journal of Approximate Reasoning*, 125, 218–239.
- CUSSENS, J. 2001. Parameter estimation in stochastic logic programs. *Mach. Learn.* 44, 3, 245–271.
- DARWICHE, A. 2004. New advances in compiling CNF into decomposable negation normal form. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain, August 22-27, 2004*, R. L. de Mántaras and L. Saitta, Eds. IOS Press, 328–332.
- DE FINETTI, B. 1931. Sul significato soggettivo della probabilità. *Fundamenta Mathematicae* 17, 298–329.
- DE RAEDT, L., FRASCONI, P., KERSTING, K. AND MUGGLETON, S., Eds. 2008. *Probabilistic Inductive Logic Programming - Theory and Applications*. Lecture Notes in Computer Science, vol. 4911. Springer.

- DE RAEDT, L., KIMMIG, A. AND TOIVONEN, H. 2007. Problog: A probabilistic prolog and its application in link discovery. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, M. M. Veloso, Ed. 2462–2467.
- DRIES, A., KIMMIG, A., MEERT, W., RENKENS, J., DEN BROECK, G. V., VLASSELAER, J. AND DE RAEDT, L. 2015. Problog2: Probabilistic logic programming. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part III*, A. Bifet, M. May, B. Zadrozny, R. Gavaldà, D. Pedreschi, F. Bonchi, J. S. Cardoso and M. Spiliopoulou, Eds. Lecture Notes in Computer Science, vol. 9286. Springer, 312–315.
- DUNG, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77, 2, 321–358.
- FIERENS, D., DEN BROECK, G. V., THON, I., GUTMANN, B. AND DE RAEDT, L. 2011. Inference in probabilistic logic programs using weighted CNF's. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, July 14-17, 2011*, F. G. Cozman and A. Pfeffer, Eds. AUAI Press, 211–220.
- GELDER, A. V., ROSS, K. A. AND SCHLIPF, J. S. 1991. The well-founded semantics for general logic programs. *J. ACM* 38, 3, 620–650.
- GELFOND, M. 2008. Answer sets. In *Handbook of Knowledge Representation*, F. van Harmelen, V. Lifschitz and B. W. Porter, Eds. Foundations of Artificial Intelligence, vol. 3. Elsevier, 285–316.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Logic Programming, Proceedings of the Fifth International Conference and Symposium, Seattle, Washington, USA, August 15-19, 1988 (2 Volumes)*, R. A. Kowalski and K. A. Bowen, Eds. MIT Press, 1070–1080.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9, 365–385.
- GUTMANN, B., THON, I. AND DE RAEDT, L. 2011. Learning the parameters of probabilistic logic programs from interpretations. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2011, Athens, Greece, September 5-9, 2011. Proceedings, Part I*, D. Gunopulos, T. Hofmann, D. Malerba and M. Vazirgiannis, Eds. Lecture Notes in Computer Science, vol. 6911. Springer, 581–596.
- HADJICHRISTODOULOU, S. AND WARREN, D. S. 2012. Probabilistic logic programming with well-founded negation. In *42nd IEEE International Symposium on Multiple-Valued Logic, ISMVL 2012, Victoria, BC, Canada, May 14-16, 2012*, D. M. Miller and V. C. Gaudet, Eds. IEEE Computer Society, 232–237.
- HUNTER, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54, 1, 47–81.
- HUNTER, A., POLBERG, S., POTYKA, N., RIENSTRA, T. AND THIMM, M. 2021. Probabilistic argumentation: A survey. *Handbook of Formal Argumentation* 2, 397–441.
- HUNTER, A., POLBERG, S. AND THIMM, M. 2018. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *CoRR abs/1802.07489*.
- HUNTER, A., POLBERG, S. AND THIMM, M. 2020. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence* 281, 103236.
- JAYNES, E. T. 1988. *The Relation of Bayesian and Maximum Entropy Methods*. Springer Netherlands, Dordrecht, 25–29.
- KOLLER, D. AND FRIEDMAN, N. 2011. Probabilistic graphical models: Principles and techniques. *MIT Press, 1231 pp., ISBN 0-262-01319-3, Knowledge Engineering Review* 26, 2, 237–238.
- LEE, J. AND WANG, Y. 2016. Weighted rules under the stable model semantics. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016*, C. Baral, J. P. Delgrande and F. Wolter, Eds. AAAI Press, 145–154.



- LI, H., OREN, N. AND NORMAN, T. J. 2011. Probabilistic argumentation frameworks. In *Theorie and Applications of Formal Argumentation - First International Workshop, TAFA 2011. Barcelona, Spain, July 16-17, 2011, Revised Selected Papers*, S. Modgil, N. Oren and F. Toni, Eds. Lecture Notes in Computer Science, vol. 7132. Springer, 1–16.
- LUKASIEWICZ, T. 1998. Probabilistic logic programming. In *13th European Conference on Artificial Intelligence, Brighton, UK, August 23-28 1998, Proceedings*, H. Prade, Ed. John Wiley and Sons, 388–392.
- MANTADELIS, T. AND BISTARELLI, S. 2020. Probabilistic abstract argumentation frameworks, a possible world view. *International Journal of Approximate Reasoning* 119, 204–219.
- MEERT, W. AND VENNEKENS, J. 2014. Inhibited effects in CP-logic. In *Probabilistic Graphical Models - 7th European Workshop, PGM 2014, Utrecht, The Netherlands, September 17-19, 2014. Proceedings*, L. C. van der Gaag and A. J. Feelders, Eds. Lecture Notes in Computer Science, vol. 8754. Springer, 350–365.
- MUGGLETON, S. *et al.* 1996. Stochastic logic programs. *Advances in inductive logic programming* 32, 254–264.
- MUISE, C. J., McILRAITH, S. A., BECK, J. C. AND HSU, E. I. 2012. Dsharp: Fast d-dnnf compilation with sharpSAT. In *Advances in Artificial Intelligence - 25th Canadian Conference on Artificial Intelligence, Canadian AI 2012, Toronto, ON, Canada, May 28-30, 2012. Proceedings*, L. Kosseim and D. Inkpen, Eds. Lecture Notes in Computer Science, vol. 7310. Springer, 356–361.
- NIELSEN, S. H. AND PARSONS, S. 2006. A generalization of dung’s abstract framework for argumentation: Arguing with sets of attacking arguments. In *Argumentation in Multi-Agent Systems, Third International Workshop, ArgMAS 2006, Hakodate, Japan, May 8, 2006, Revised Selected and Invited Papers*, N. Maudet, S. Parsons and I. Rahwan, Eds. Lecture Notes in Computer Science, vol. 4766. Springer, 54–73.
- NILSSON, N. 2014. *Understanding Beliefs*. The MIT Press Essential Knowledge Series. MIT Press.
- PEARL, J. 1989. *Probabilistic Reasoning in Intelligent Systems - Networks of Plausible Inference*. Morgan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann.
- POLBERG, S., HUNTER, A. AND THIMM, M. 2017. Belief in attacks in epistemic probabilistic argumentation. In *Scalable Uncertainty Management - 11th International Conference, SUM 2017, Granada, Spain, October 4-6, 2017, Proceedings*, S. Moral, O. Pivert, D. Sánchez and N. Marín, Eds. Lecture Notes in Computer Science, vol. 10564. Springer, 223–236.
- POOLE, D. 2008. The independent choice logic and beyond. See [De Raedt et al. \(2008\)](#), 222–243.
- RIGUZZI, F. 2015. The distribution semantics is well-defined for all normal programs. In *Proceedings of the 2nd International Workshop on Probabilistic Logic Programming co-located with 31st International Conference on Logic Programming (ICLP 2015), Cork, Ireland, August 31st, 2015*, F. Riguzzi and J. Vennekens, Eds. CEUR Workshop Proceedings, vol. 1413. CEUR-WS.org, 69–84.
- ROCHA, V. H. N. AND COZMAN, F. G. 2022. A credal least undefined stable semantics for probabilistic logic programs and probabilistic argumentation. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31 – August 5, 2022*, G. Kern-Isberner, G. Lakemeyer and T. Meyer, Eds.
- SACCÀ, D. AND ZANIOLO, C. 1990. Stable models and non-determinism in logic programs with negation. In *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, April 2-4, 1990, Nashville, Tennessee, USA*, D. J. Rosenkrantz and Y. Sagiv, Eds. ACM Press, 205–217.
- SATO, T. 1995. 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*, L. Sterling, Ed. MIT Press, 715–729.
- SATO, T. AND KAMEYA, Y. 2008. New advances in logic-based probabilistic modeling by PRISM. See [De Raedt et al. \(2008\)](#), 118–155.

- STEDE, M., AFANTENOS, S. D., PELDSZUS, A., ASHER, N. AND PERRET, J. 2016. Parallel discourse annotations on a corpus of short texts. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*, N. Calzolari, K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk and S. Piperidis, Eds. European Language Resources Association (ELRA).
- VENNEKENS, J. 2013. Negation in the head of cp-logic rules. *CoRR abs/1312.6156*.
- VENNEKENS, J., DENECKER, M. AND BRUYNOOGHE, M. 2009. CP-logic: A language of causal probabilistic events and its relation to logic programming. *TPLP* 9, 3, 245–308.
- WU, Y., CAMINADA, M. AND GABBAY, D. M. 2009. Complete extensions in argumentation coincide with 3-valued stable models in logic programming. *Studia Logica* 93, 2–3, 383–403.