

The Complexity of Inferences and Explanations in Probabilistic Logic Programming

Fabio G. Cozman¹ and Denis D. Mauá²

¹ Escola Politécnica, Universidade de São Paulo - Brazil

² Instituto de Matemática e Estatística, Universidade de São Paulo - Brazil

Abstract. A popular family of probabilistic logic programming languages combines logic programs with independent probabilistic facts. We study the complexity of marginal inference, most probable explanations, and maximum a posteriori calculations for propositional/relational probabilistic logic programs that are acyclic/definite/stratified/normal/disjunctive. We show that complexity classes Σ_k and PP^{Σ_k} (for various values of k) and NP^{PP} are all reached by such computations.

1 Introduction

The goal of this paper is to shed light on the computational complexity of inference for probabilistic logic programs interpreted in the spirit of Sato's distribution semantics [25]; that is, we have logic programs where some facts are annotated with probabilities, so as to define probability distributions over models. This framework has been shown to be quite useful in modeling practical problems [15, 24].

The distribution defined by a probabilistic logic program can be used to answer many queries of interest. Two common queries are to compute the probability of some ground atom given evidence (inference), and to find a (partial) interpretation that maximizes probability while being consistent with evidence (MPE/MAP).

We present results on the complexity of *acyclic*, *definite*, *stratified*, *normal* and *disjunctive* probabilistic logic programs; these results are summarized by Table 1. While most semantics agree on stratified programs, there is less consensus on non-stratified programs. Here we examine two semantics: the *credal semantics*, based on stable models, and the *well-founded semantics*.

We start in Sections 2 and 3 by reviewing relevant background on probabilistic logic programs and on complexity theory. Our contributions appear in Section 4. These results are further discussed in the concluding Section 5.

2 Background

The results in this paper depend on an understanding of logic and answer set programming; the topic is dense and cannot be described in detail in the space

	Propositional			Bounded arity		
	Inferential	MPE	MAP	Inferential	MPE	MAP
Acyclic normal	PP	NP	NP ^{PP}	PP ^{NP}	Σ_2^P	NP ^{PP}
Definite (positive query)	PP	NP	NP ^{PP}	PP ^{NP}	Σ_2^P	NP ^{PP}
Stratified normal	PP	NP	NP ^{PP}	PP ^{NP}	Σ_2^P	NP ^{PP}
Normal, credal	PP ^{NP}	Σ_2^P	NP ^{PP}	PP ^{Σ_2^P}	Σ_3^P	NP ^{PP}
Normal, well-founded	PP	NP	NP ^{PP}	PP ^{NP}	Σ_2^P	NP ^{PP}
Disjunctive, credal	PP ^{Σ_2^P}	Σ_3^P	NP ^{PP}	PP ^{Σ_3^P}	Σ_4^P	NP ^{PP}

Table 1. Summary of complexity results presented in this paper (all entries indicate completeness with respect to many-one reductions). Entries containing known results have orange background (grey background if printed in black-and-white) [6–8].

we have here. We just mention the main concepts, and refer the reader to any in-depth presentation in the literature [9, 14].

We have a vocabulary consisting of *predicates*, *constants*, and *logical variables*; a *term* is a constant or logical variable, and an *atom* is a predicate of *arity* k associated with k terms; a *ground atom* is an atom without logical variables. We often resort to *grounding* to produce ground atoms. A *disjunctive logic program* (DLP) consists of a set of rules written as

$$A_1 \vee \dots \vee A_h :- B_1, \dots, B_{b'}, \mathbf{not} B_{b'+1}, \dots, \mathbf{not} B_b .$$

where each A_i and B_i is an atom. The lefthand side is the *head* of the rule; the remainder is its *body*. A rule without disjunction (i.e., $h = 1$) and with empty body, written simply as $A_1 .$, is a *fact*. A program without disjunction is a *normal logic program*; a normal logic program without negation is a *definite program*; finally, a program without variables is a *propositional program*. The *dependency graph* of a program is the graph where each atom is a vertex and there are arcs from atoms in the bodies to atoms in the heads of a same rule; a normal logic program is *acyclic* when the dependency graph of its grounding is acyclic. A normal logic program is (locally) *stratified* when the dependency graph of its grounding has no cycles containing an arc involving a negated literal.

The *Herbrand base* of a program is the set of all ground atoms built from constants and predicates in the program. An *interpretation* is a set of ground literals mentioning exactly once each atom in the Herbrand base. A *model* is an interpretation that *satisfies* every grounding of a rule (a rule is satisfied iff the interpretation contains all of $B_1, \dots, B_{b'}$, none of $B_{b'+1}, \dots, B_b$, and some of A_1, \dots, A_h). A *minimal model* minimizes the number of non-negated literals.

The most common semantics for DLPs is the *stable model* semantics. Given a program \mathbf{P} and an interpretation \mathcal{I} , define their *reduct* to be program obtained by removing from \mathbf{P} every rule whose body is not satisfied by \mathcal{I} . An interpretation is a *stable model* if it is a minimal model of its reduct. A normal program may have zero, one or several stable models. *Brave reasoning* asks whether there is

a stable model containing a specific literal (possibly returning one). *Cautious reasoning* asks whether a specific literal appears in all stable models (possibly listing all).

Here is an example: Several robots can perform, each, one of three operations, called “red”, “green”, “yellow”. A robot placed in a site also covers adjacent sites, so there is no need to place same-color robots on adjacent sites. There is a list of robots and a list of (one-way) connections between sites; the goal is to distribute the robots and verify whether the sites are connected. This disguised 3-coloring problem can be encoded as [14]:

$$\begin{aligned} & \text{color}(X, \text{red}) \vee \text{color}(X, \text{green}) \vee \text{color}(X, \text{yellow}) :- \text{site}(X). \\ & \text{clash} :- \text{not clash}, \text{edge}(X, Y), \text{color}(X, C), \text{color}(Y, C). \\ & \text{path}(X, Y) :- \text{edge}(X, Y). \quad \text{path}(X, Y) :- \text{edge}(X, Z), \text{path}(Z, Y). \end{aligned}$$

We might have a database of facts, consisting of a list of sites and their connections, say $\text{site}(s_1), \text{site}(s_2), \dots, \text{edge}(s_1, s_4)$, and so on. Each stable model of this program is a possible placement (a 3-coloring) and a list of paths between sites.

An alternative semantics for normal logic programs is the *well-founded semantics*; a model under this semantics might fix only a truth-value for some of the atoms (leaving the remaining atoms undefined) [31]. One way to define the well-founded semantics is as follows [4]. Write $\text{LFT}_{\mathbf{P}}(\mathcal{I})$ to mean the least fixpoint of $\mathbb{T}_{\mathbf{P}z}$, where $\mathbb{T}_{\mathbf{P}}$ is a transformation such that: atom A is in $\mathbb{T}_{\mathbf{P}}(\mathcal{I})$ iff there is grounded rule with head A with the whole body true in interpretation \mathcal{I} . Then the well-founded semantics of \mathbf{P} consists of those atoms A that are in the least fixpoint of $\text{LFT}_{\mathbf{P}}(\text{LFT}_{\mathbf{P}}(\cdot))$ plus the literals $\neg A$ for those atoms A that are *not* in the greatest fixpoint of $\text{LFT}_{\mathbf{P}}(\text{LFT}_{\mathbf{P}}(\cdot))$.

We also need standard concepts from complexity theory: languages (sets of strings), decision problems (deciding whether input is in language), complexity classes (sets of languages), many-one reductions [22]. We use well-known complexity classes such as P, NP, PP. We also consider oracle machines and corresponding complexity classes such as Π_i^P and Σ_i^P (the so-called *polynomial hierarchy*). We also use Wagner’s *polynomial counting hierarchy* defined as the smallest set of classes containing P and, recursively, for any class C in the polynomial counting hierarchy, the classes PP^C , NP^C , and coNP^C [29, 32].

3 Probabilistic logic programming

In this paper we focus on a particularly simple combination of logic programming and probabilities [23, 25]. A *probabilistic disjunctive logic program*, abbreviated PDLP, is a pair $\langle \mathbf{P}, \mathbf{PF} \rangle$ consisting of a disjunctive logic program \mathbf{P} and a set of *probabilistic facts* \mathbf{PF} . A probabilistic fact is a pair consisting of an atom A and a probability value α , written as $\alpha :: A$. [15]. Note that we allow a probabilistic fact to contain logical variables. As an instance of PDLP, take our running example on robots and sites: to generate a random graph over a set of five sites add the rules: $0.5 :: \text{edge}(X, Y).$, $\text{site}(s_1).$, $\text{site}(s_2).$, $\text{site}(s_3).$, $\text{site}(s_4).$, $\text{site}(s_5).$. If \mathbf{P} is a normal logic program, we just write *probabilistic logic program*,

abbreviated PLP. If \mathbf{P} is normal and acyclic/definite/stratified, we say the PLP is acyclic/definite/stratified.

To build the semantics of a PDLP, we first take its grounding. From a PDLP with n ground probabilistic facts, we can generate 2^n DLPS: for each probabilistic fact $\alpha :: A.$, either keep fact $A.$ with probability α , or erase $A.$ with probability $1 - \alpha$. A *total choice* is a subset of the set of ground probabilistic facts that is selected to be kept (other grounded probabilistic facts are discarded). For any total choice θ we obtain a DLP $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ with probability $\prod_{A_i \in \theta} \alpha_i \prod_{A_i \notin \theta} (1 - \alpha_i)$. The distribution over total choices induces a distribution over DLPS.

We first define a semantics proposed by Lukasiewicz [18, 19]. A *probability model* for a PDLP $\langle \mathbf{P}, \mathbf{PF} \rangle$ is a probability measure \mathbb{P} over interpretations, such that (i) every interpretation \mathcal{I} with $\mathbb{P}(\mathcal{I}) > 0$ is a stable model of $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ for the total choice θ that agrees with \mathcal{I} on the probabilistic facts; and (ii) the probability of a total choice θ is $\mathbb{P}(\theta) = \prod_{A_i \in \theta} \alpha_i \prod_{A_i \notin \theta} (1 - \alpha_i)$. The set of all probability models for a PDLP is the semantics of the program; note that if a PDLP does not have stable models for some total choice, there is no semantics for it (the program is *inconsistent*). Because a set of probability measures is often called a *credal set* [2]; we adopt the term *credal semantics*.

If \mathbf{P} is definite, then $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ is definite for any θ , and $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ has a unique minimal model that is also its unique stable/well-founded model. Thus the distribution over total choices induces a single probability model. This is Sato's *distribution semantics* [25]. Similarly, suppose that \mathbf{P} is acyclic or stratified; then $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ is respectively acyclic or stratified for any θ , and $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ has a unique stable model that is also its unique well-founded model [1].

Given a consistent PDLP whose credal semantics is the credal set \mathbb{K} , we may be interested in computing *lower conditional probabilities*, defined as $\underline{\mathbb{P}}(\mathbf{Q}|\mathbf{E}) = \inf_{\mathbb{P} \in \mathbb{K}: \mathbb{P}(\mathbf{E}) > 0} \mathbb{P}(\mathbf{Q}|\mathbf{E})$ or *upper conditional probabilities*, defined as $\overline{\mathbb{P}}(\mathbf{Q}|\mathbf{E}) = \sup_{\mathbb{P} \in \mathbb{K}: \mathbb{P}(\mathbf{E}) > 0} \mathbb{P}(\mathbf{Q}|\mathbf{E})$, where \mathbf{Q} and \mathbf{E} are consistent set of literals. Note that we leave conditional lower/upper probabilities undefined when $\overline{\mathbb{P}}(\mathbf{E}) = 0$ (that is, when $\mathbb{P}(\mathbf{E}) = 0$ for every probability model).

Consider again our running example. Suppose we have a graph over five sites, with edges (to save space, e means edge):

$0.5 :: e(s_4, s_5). \quad e(s_1, s_3). \quad e(s_1, s_4). \quad e(s_2, s_1). \quad e(s_2, s_4). \quad e(s_3, s_5). \quad e(s_4, s_3).$

That is, we have an edge $e(s_4, s_5)$ which appears with probability 0.5. If this edge is kept, there are 6 stable models; if it is discarded, there are 12 stable models. If additional facts $\text{color}(s_2, \text{red}).$ and $\text{color}(s_5, \text{green}).$ are given, then there is a single stable model if $\text{edge}(s_4, s_5)$ is kept, and 2 stable models if it is discarded. Then, $\underline{\mathbb{P}}(\text{color}(s_4, \text{green})) = 0$ and $\overline{\mathbb{P}}(\text{color}(s_4, \text{green})) = 1/2$.

A different semantics is defined by Hadjichristodoulou and Warren [17] for *normal* PLPs: they allow probabilities directly over well-founded models, thus allowing probabilities over atoms that are *undefined*. That is, given a PLP $\langle \mathbf{P}, \mathbf{PF} \rangle$, associate to each total choice θ the unique well-founded model of $\mathbf{P} \cup \mathbf{PF}^{\downarrow\theta}$ to θ ; the unique distribution over total choices induces a unique distribution over well-founded models. To conclude, we note that one can find other semantics in the literature that serve deserve further study [3, 5, 18, 21, 26].

4 Complexity results

We consider three different problems in this paper: (marginal) inference, most probable explanation (MPE), and maximum a posteriori (MAP).

In the following problem definitions a PDLP $\langle \mathbf{P}, \mathbf{PF} \rangle$ is always specified using rational numbers as probability values, and with a bound on the arity of predicates (so the Herbrand base is always polynomial in the input size). A query (\mathbf{Q}, \mathbf{E}) is always a pair of sets of consistent literals (consistent here means that the set does not contain both a literal and its negation). The set \mathbf{E} is called *evidence*. The symbol \mathbf{M} denotes a set of atoms in the Herbrand base of the union of the program \mathbf{P} and all the facts in \mathbf{PF} . The symbol γ is always a rational number in $[0, 1]$.

The *inferential complexity* of a class of PDLPs is the complexity of the following decision problem: with **input** equal to a PDLP $\langle \mathbf{P}, \mathbf{PF} \rangle$, a query (\mathbf{Q}, \mathbf{E}) , and a number γ , the **output** is whether or not $\mathbb{P}(\mathbf{Q}|\mathbf{E}) > \gamma$; by convention, the input is rejected if $\mathbb{P}(\mathbf{E}) = 0$.

The *MPE complexity* of a class of PDLPs is the complexity of the following decision problem: with **input** equal to a PDLP $\langle \mathbf{P}, \mathbf{PF} \rangle$, evidence \mathbf{E} , and a number γ , the **output** is whether or not there is an interpretation \mathcal{I} that agrees with \mathbf{E} and satisfies $\mathbb{P}(\mathcal{I}) > \gamma$.

The *MAP complexity* of a class of PDLPs is the complexity of the following decision problem: with **input** equal to a PDLP $\langle \mathbf{P}, \mathbf{PF} \rangle$, a set \mathbf{M} , and a number γ , the **output** is whether or not there is a consistent set of literals \mathbf{Q} mentioning all atoms in \mathbf{M} such that $\mathbb{P}(\mathbf{Q}|\mathbf{E}) > \gamma$; by convention, the input is rejected if $\mathbb{P}(\mathbf{E}) = 0$.

The contributions of this paper are summarized by Table 1. Darker entries are already known [6–8], and some entries on acyclic PLPs can also be found in work by Ceylan et al. [5]. All entries in this table indicate completeness with respect to many-one reductions. In this section we prove these facts through a series of results.³ In all proofs the argument for membership depends on the complexity of logical reasoning on logic programs that are obtained by fixing all total choices; this suffices to even make decisions concerning conditional probabilities (using for instance techniques by Park [10, Theorem 11.5]).

Theorem 1 *The MPE complexity of acyclic propositional PLPs is NP-hard, and of stratified propositional PLPs is in NP. The MPE complexity of acyclic PLPs is Σ_2^P -hard, and of stratified PLPs is in Σ_2^P .*

Proof. Membership for stratified propositional PLPs: “guess” a polynomial-sized interpretation consistent with evidence, and then decide if its probability exceeds a given threshold in polynomial time (by checking the stable model [12, Table 4]). Hardness: use acyclic propositional PLPs to encode Bayesian networks [11].

³ IMPORTANT NOTE: Due to space restrictions; we only present proof sketches; the reader can find the complete proofs at <http://sites.poli.usp.br/p/fabio.cozman/Publications/Article/cozman-maua-ecsqaru2017.pdf>.

To prove membership for stratified PLPs, note that deciding whether a given interpretation is a stable model of a stratified logic program can be reduced to an instance of cautious reasoning. Hence, we can “guess” an interpretation consistent with evidence, then decide whether it is a stable model using a $\mathbf{P}^{\mathbf{NP}}$ oracle [12, Table 5]. To obtain Σ_2^P -hardness, use an encoding employed by Eiter et al. [12]. Suppose we have formula $\phi = \exists \mathbf{X} : \neg \exists \mathbf{Y} : \varphi(\mathbf{X}, \mathbf{Y})$, where $\varphi(\mathbf{X}, \mathbf{Y})$ is a propositional formula in 3CNF with sets \mathbf{X} and \mathbf{Y} of propositional variables. Deciding satisfiability of such formulas is a Σ_2^P -complete problem [27]. Introduce $0.5 :: \mathbf{x}$. with predicate x for each x in \mathbf{X} . A clause c in φ contains $k \in \{0, \dots, 3\}$ propositional variables from \mathbf{Y} . Introduce a predicate c of arity k , and for each predicate c introduce a set of rules: For each one of the 2^k groundings \mathbf{y}' of the logical variables \mathbf{Y}' in c , if \mathbf{y}' satisfies c (for all assignments of \mathbf{X}), introduce a fact $c(\mathbf{y}')$; if \mathbf{y}' does not satisfy c (for some assignment of \mathbf{X}), introduce $3 - k$ rules of the form $c(\mathbf{y}') :- [\mathbf{not}] \mathbf{x}$., where **not** appears in the rule depending on whether x is preceded by negation or not in the clause c . The formula φ is encoded by the rule $\mathbf{cnf} :- \mathbf{c}_1, \mathbf{c}_2, \dots$. (conjunction extends over all clauses). Then the MPE with evidence $\{\neg \mathbf{cnf}\}$ and threshold $\gamma = 0$ decides whether ϕ is satisfiable. \square

As definite programs are stratified, they are already covered by previous results. However, it makes sense to assume that any query with respect to such a program will also be *positive* in the sense that it only contains non-negated literals. Even then we have the same complexity as stratified programs:

Theorem 2 *Assume all queries are positive. The inferential complexity of definite propositional PLPs is PP-complete, and of definite PLPs is $\mathbf{PP}^{\mathbf{NP}}$ -complete. The MPE complexity of definite propositional PLPs is NP-complete, and of definite PLPs is Σ_2^P -complete.*

Proof. Membership follows from results for stratified PLPs [7]. To show hardness for propositional programs, consider a 3CNF formula φ over variables x_1, \dots, x_n , and obtain a new monotone formula $\tilde{\varphi}$ by replacing every literal $\neg x_i$ by a fresh variable y_i . Now the formula φ has M satisfying assignments iff the formula $\tilde{\varphi} \wedge (\bigwedge_i x_i \vee y_i) \vee \bigvee_i (x_i \wedge y_i)$ has $M + 2^{2n} - 3^n$ satisfying assignments [16, Proposition 4]. The latter formula is monotone (i.e., contains no negated variables), so we can encode it as a definite program using probabilistic facts $0.5 :: \mathbf{x}$. to represent each logical variable, \mathbf{c}_i to represent clauses, and \mathbf{cnf} to represent the value of the formula. To decide whether the number of solutions of φ exceeds M , verify whether $\mathbb{P}(\mathbf{cnf}) > (2^{2n} - 3^n + M)/2^{2n}$. To decide if there is a solution, decide the MPE with evidence $\{\mathbf{cnf}\}$ and threshold $2^{2n} - 3^n$. The same reasoning applies to definite PLPs, by building existential quantification over part of the variables as in the proof of Theorem 1. \square

Non-stratified programs climb one step up in the polynomial hierarchy:

Theorem 3 *Assume the credal semantics for PLPs. The MPE complexity of propositional PLPs is Σ_2^P -complete, and of PLPs is Σ_3^P -complete.*

Proof. Adapt the proof of Theorem 1, as follows. For the propositional case, membership obtains as cautious reasoning is coNP -complete [12, Table 2], and Σ_2^P -hardness obtains by encoding a formula $\phi = \exists \mathbf{X} : \forall \mathbf{Y} : \varphi(\mathbf{X}, \mathbf{Y})$, where $\varphi(\mathbf{X}, \mathbf{Y})$ is a propositional formula in 3DNF. Encode each $x \in \mathbf{X}$ as before, and introduce y_i and ny_i for each y_i in \mathbf{Y} , together with rules $y_i :- \text{not ny}_i.$ and $\text{ny}_i :- \text{not } y_i.$. Then encode φ by rules $\text{dnf} := d_j$ (where each d_j is a conjunct of φ); the MPE of this program with evidence $\{\text{dnf}\}$ and threshold 0 decides whether ϕ is satisfiable (the “inner” universal quantifier is “produced” by going over all stable models when doing cautious reasoning). For the arity-bounded case, membership obtains as cautious reasoning is Π_2^P -complete [12, Table 5], and Σ_3^P -hardness obtains by a combination of strategies used in the proof of Theorem 1 and in the proof for propositional PLPs. That is, encode a formula $\phi = \exists \mathbf{X} : \forall \mathbf{Y} : \exists \mathbf{Z} : \varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, where φ is in CNF. \square

Now consider inferential complexity under the credal semantics, for PDLPs. The credal semantics of a PDLP is a credal set that dominates an infinite monotone Choquet capacity [6]. This result is important because it implies that $\underline{\mathbb{P}}(\mathcal{M}) = \sum_{\theta \in \Theta: \Gamma(\theta) \subseteq \mathcal{M}} \mathbb{P}(\theta)$ and $\overline{\mathbb{P}}(\mathcal{M}) = \sum_{\theta \in \Theta: \Gamma(\theta) \cap \mathcal{M} \neq \emptyset} \mathbb{P}(\theta)$, where Θ is the set of total choices and Γ maps a total choice to the set of resulting stable models. Also, we have that $\underline{\mathbb{P}}(A|B) = \underline{\mathbb{P}}(A \cap B) / (\underline{\mathbb{P}}(A \cap B) + \overline{\mathbb{P}}(A^c \cap B))$ whenever $\underline{\mathbb{P}}(A \cap B) + \overline{\mathbb{P}}(A^c \cap B) > 0$; otherwise, either $\underline{\mathbb{P}}(A|B) = 1$ when $\underline{\mathbb{P}}(A \cap B) + \overline{\mathbb{P}}(A^c \cap B) = 0$ and $\overline{\mathbb{P}}(A \cap B) > 0$, or $\underline{\mathbb{P}}(A|B)$ is undefined. Similarly, $\overline{\mathbb{P}}(A|B) = \overline{\mathbb{P}}(A \cap B) / (\overline{\mathbb{P}}(A \cap B) + \underline{\mathbb{P}}(A^c \cap B))$, when $\overline{\mathbb{P}}(A \cap B) + \underline{\mathbb{P}}(A^c \cap B) > 0$, with similar special cases. Using these results we see that computing lower and upper probabilities can be reduced to going through the total choices and running brave/cautious inference for each total choice [6].

Theorem 4 *Assume the credal semantics for PDLPs. The inferential complexity of propositional PDLPs is $\text{PP}^{\Sigma_2^P}$ -complete, and of PDLPs is $\text{PP}^{\Sigma_3^P}$ -complete.*

Proof. To prove membership for propositional PDLPs, note that once a total choice is guessed, the cost of checking whether a set of literals holds under cautious reasoning is in Π_2^P [12, Table 2]. To obtain $\text{PP}^{\Sigma_2^P}$ -hardness, consider a formula $\phi(\mathbf{X}) = \forall \mathbf{Y} : \neg \forall \mathbf{Z} : \varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, where $\varphi(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is a propositional formula in 3DNF with conjuncts d_j and sets of propositional variables \mathbf{X} , \mathbf{Y} , and \mathbf{Z} . Deciding whether the number of truth assignments to \mathbf{X} that satisfy the formula is strictly larger than an integer M is a $\text{PP}^{\Sigma_2^P}$ -complete problem [32, Theorem 7]. To emulate counting, introduce a predicate x_i for each propositional variable x_i in \mathbf{X} , associated with a probabilistic fact $0.5 :: x_i.$. And to encode $\phi(\mathbf{X})$, we combine the use of stable sets as in the proof of Theorem 3 with an adapted version of a proof by Eiter and Gottlob [13] on disjunctive programming. We focus on the latter construction here. Introduce predicates z_i and nz_i for each propositional variable z_i in \mathbf{Z} , and auxiliary predicate w , together with the rules $z_i \vee \text{nz}_i., z_i :- w.,$ and $\text{nz}_i :- w.$ for each z_i in \mathbf{Z} , plus the rule $w :- L_j^1, L_j^2, L_j^3.$ for each conjunct d_j , where L_j^r is obtained from L , the r th literal of d_j , as follows: (1) if $L = z_i$, then $L_j^r = z_i$; (2) if $L = \neg z_i$, then $L_j^r = \text{nz}_i$; (3) if $L = x_i$, then

$L_j^r = x_i$; (4) if $L = \neg x_i$, then $L_j^r = \mathbf{not} x_i$; (5) if $L = y_i$, then $L_j^r = y_i$; (6) if $L = \neg y_i$, then $L_j^r = \mathbf{not} y_i$. Finally, introduce $\mathbf{nw} :- \mathbf{not} \mathbf{w}.$ Now reason as follows. To decide whether $\mathbb{P}(\mathbf{nw} = \mathbf{true}) > \gamma$, we must go through all possible total choices; each one of them has probability 2^{-n} where n is the length of \mathbf{X} . For each total choice, we must run cautious inference; this is done by going through all stable models, and verifying whether \mathbf{nw} is \mathbf{true} in all of them. For each truth assignment of \mathbf{Y} , the program has a stable model where \mathbf{w} is \mathbf{true} iff for all truth assignments of \mathbf{Z} we have that φ holds [13, Theorem 3.2]. Hence for fixed \mathbf{y} , resulting stable models have \mathbf{nw} as \mathbf{true} iff $\forall \mathbf{Z} : \varphi(\mathbf{x}, \mathbf{y}, \mathbf{Z})$ is false (\mathbf{x} is fixed by the selected total choice). Thus if we take $\gamma = M/2^{-n}$, we obtain that $\mathbb{P}(\mathbf{nw} = \mathbf{true}) > \gamma$ decides whether $\phi(\mathbf{X})$ has a number of satisfying truth assignments of \mathbf{X} that is strictly larger than M .

To prove membership for PDLPs, note that once a total choice is guessed, the cost of checking whether a set of literals holds under cautious reasoning is in Π_3^P [12, Table 5]. To obtain $\text{PP}^{\Sigma_3^P}$ -hardness, we use a combination of strategies used in the proof of Theorem 1 and in the proof for propositional PLPs (previous paragraph) to encode a formula $\phi(\mathbf{X}) = \forall \mathbf{Y} : \neg \forall \mathbf{Z} : \exists \mathbf{V} : \varphi(\mathbf{V}, \mathbf{X}, \mathbf{Y}, \mathbf{Z})$. \square

Theorem 5 *Assume the credal semantics for PDLPs. The MPE complexity of propositional PDLPs is Σ_3^P -complete, and of PDLPs is Σ_4^P -complete.*

Proof. Membership and hardness are proved by adapting arguments in the proofs of Theorems 3 and 4, using the complexity of cautious reasoning [12, Tables 2 and 5], and encodings of DNF and CNF formulas as before. \square

Theorem 6 *Assume the well-founded semantics for PLPs. The MPE complexity of propositional PLPs is NP-complete, and of PLPs is Σ_2^P -complete.*

Proof. Hardness follows from Theorem 1 (in both cases). Membership: use the argument in the proof of Theorem 1, employing complexity of logical inference of propositional case [9] and the bounded-arity case [8] as appropriate. \square

Theorem 7 *Assume the credal semantics both for PLPs and for PDLPs. The MAP complexity of propositional PLPs is NP^{PP} -hard, and of PDLPs is in NP^{PP} .*

Proof. Hardness: a propositional PLP can encode a Bayesian network with binary variables, and MAP in such networks is NP^{PP} -complete [11]. To show membership, reason in two steps: one can solve MAP by first guessing literals for the MAP-predicates that are not fixed by evidence, and then running inference in an $\text{PP}^{\Sigma_3^P}$ oracle. That is, the decision problem is in $\text{NP}^{\text{PP}^{\Sigma_3^P}}$. Now resort to a theorem by Toda and Watanabe [28] that shows that, for any k , $\text{PPP}^{\Sigma_k^P}$ collapses to P^{PP} , to note that the decision problem of interest is in NP^{PP} . \square

Theorem 8 *Assume the credal semantics for PLPs and positive queries. The MAP complexity of definite propositional PLPs is NP^{PP} -hard, and of PLPs is in NP^{PP} .*

Proof. Membership: Theorem 7. Hardness: MAP for Bayesian networks described without negation is already NP^{PP} -hard [20, Theorem 5]. \square

5 Conclusion

As conveyed by Table 1, we have presented several novel results concerning the complexity of probabilistic logic programming, both analyzing conditional probabilities (inferences) and explanations (MPE and MAP) — we note that previous work has not examined these latter problems in the context of probabilistic logic programming. These computations go up several layers within the counting hierarchy, reaching some interesting complexity classes that are rarely visited. In particular, MAP is *always* NP^{PP} -complete, a rather interesting result.

A future step is to obtain the complexity of relational programs without bounds on arity (exponential complexity is sure to appear), and perhaps the complexity of programs with functions (with suitable restrictions to guarantee decidability). The complexity of other constructs, such as aggregates, should also be explored. Future work should look at “query” complexity; that is, the complexity of computing inferences when the program is fixed and the query varies — this is akin to data complexity as studied in database theory [7].

Acknowledgements

The first author is partially supported by CNPq, grant 308433/2014-9. The second author received support from the São Paulo Research Foundation (FAPESP), grant 2016/01055-1. The work reported in this paper was partially funded by FAPESP grant #2015/21880-4 (project Proverbs).

References

1. K. R. Apt and M. Bezem. Acyclic programs. *New Generation Computing*, 9:335–363, 1991.
2. T. Augustin, F. P. A. Coolen, G. de Cooman, and M. C. M. Troffaes. *Introduction to Imprecise Probabilities*. Wiley, 2014.
3. C. Baral, M. Gelfond, and N. Rushton. Probabilistic reasoning with answer sets. *Theory and Practice of Logic Programming*, 9(1):57–144, 2009.
4. C. Baral and V. Subrahmanian. Dualities between alternative semantics for logic programming and nonmonotonic reasoning. *J. of Automated Reasoning*, 10(3):399–420, 1993.
5. Í. Í. Ceylan, T. Lukasiewicz, and R. Peñaloza. Complexity results for probabilistic Datalog[±]. In *European Conf. on Artificial Intelligence*, pages 1414–1422, 2016.
6. F. G. Cozman and D. D. Mauá. The structure and complexity of credal semantics. In *Workshop on Probabilistic Logic Programming*, pages 3–14, 2016.
7. F. G. Cozman and D. D. Mauá. Probabilistic graphical models specified by probabilistic logic programs: Semantics and complexity. In *Conf. on Probabilistic Graphical Models — JMLR Proceedings*, volume 52, pages 110–121, 2016.
8. F. G. Cozman and D. D. Mauá. The well-founded semantics of cyclic probabilistic logic programs: meaning and complexity. In *Encontro Nacional de Inteligência Artificial e Computacional*, pages 1–12, 2016.
9. E. Dantsin, T. Eiter, and A. Voronkov. Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3):374–425, 2001.
10. A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge, 2009.

11. C. Polpo de Campos and F. G. Cozman. The inferential complexity of Bayesian and credal networks. In *IJCAI*, pages 1313–1318, 2005.
12. T. Eiter, W. Faber, M. Fink, and S. Woltran. Complexity results for answer set programming with bounded predicate arities and implications. *Annals of Mathematics and Artificial Intelligence*, 5:123–165, 2007.
13. T. Eiter and G. Gottlob. On the computational cost of disjunctive logic programming: Propositional case. *Annals of Mathematics and Artificial Intelligence*, 15:289–323, 1995.
14. T. Eiter, G. Ianni, and T. Krennwalner. Answer set programming: a primer. In *Reasoning Web*, pages 40–110. Springer-Verlag, 2009.
15. D. Fierens, G. Van den Broeck, J. Renkens, D. Shrerionov, B. Gutmann, G. Janssens, and L. de Raedt. Inference and learning in probabilistic logic programs using weighted Boolean formulas. *Theory and Practice of Logic Programming*, 15(3):358–401, 2014.
16. J. Goldsmith, M. Hagen, and M. Mundhenk. Complexity of DNF minimization and isomorphism testing for monotone formulas. *Information and Computation*, 206(6):760–775, 2008.
17. S. Hadjichristodoulou and D. S. Warren. Probabilistic logic programming with well-founded negation. *Int. Symp. on Multiple-Valued Logic*, pages 232–237, 2012.
18. T. Lukasiewicz. Probabilistic description logic programs. In *Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 737–749, 2005.
19. T. Lukasiewicz. Probabilistic description logic programs. *In. J. of Approximate Reasoning*, 45(2):288–307, 2007.
20. D. D. Mauá, C. Polpo de Campos, F. G. Cozman. The complexity of MAP inference in Bayesian networks specified through logical languages. *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 889–895, 2015.
21. S. Michels, A. Hommersom, P. J. F. Lucas, and M. Velikova. A new probabilistic constraint logic programming language based on a generalised distribution semantics. *Artificial Intelligence J.*, 228:1–44, 2015.
22. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
23. D. Poole. Probabilistic Horn abduction and Bayesian networks. *Artificial Intelligence*, 64:81–129, 1993.
24. D. Poole. The Independent Choice Logic and beyond. In L. De Raedt, P. Frasconi, K. Kersting, and S. Muggleton, editors, *Probabilistic Inductive Logic Programming*, volume 4911 of *Lecture Notes in Computer Science*, pages 222–243. Springer, 2008.
25. T. Sato. A statistical learning method for logic programs with distribution semantics. In *Int. Conf. on Logic Programming*, pages 715–729, 1995.
26. T. Sato, Y. Kameya, and N.-F. Zhou. Generative modeling with failure in PRISM. In *Int. Joint Conf. on Artificial Intelligence*, pages 847–852, 2005.
27. L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3:1, 1–22, 1976.
28. S. Toda and O. Watanabe. Polynomial-time 1-Turing reductions from #PH to #P. *Theoretical Computer Science*, 100:205–221, 1992.
29. J. Tóran. Complexity classes defined by counting quantifiers. *J. of the ACM*, 38(3):753–774, 1991.
30. L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.
31. A. van Gelder, J. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *J. Association for Computing Machinery*, 38(3):620–650, 1991.
32. K. W. Wagner. The complexity of combinatorial problems with succinct input representation. *Acta Informatica*, 23:325–356, 1986.