

Deep Probabilistic Logic Programming

Arnaud Nguembang Fadja Evelina Lamma Fabrizio Riguzzi

Dipartimento di Ingegneria – University of Ferrara

Dipartimento di Matematica e Informatica – University of Ferrara

[arnaud.nguembangfadja,evelina.lamma,fabrizio.riguzzi]@unife.it



Introduction

- Probabilistic logic programming is a powerful tool for reasoning with uncertain relational models
- Learning probabilistic logic programs is expensive due to the high cost of inference.
- We consider a restriction of the language of Logic Programs with Annotated Disjunctions called **hierarchical PLP** in which clauses and predicates are hierarchically organized.
- Hierarchical PLP is truth-functional and equivalent to the product fuzzy logic.
- Inference then is much cheaper as a simple dynamic programming algorithm similar to PRISM is sufficient

Probabilistic Logic Programming

- Distribution Semantics [Sato ICLP95]
- A probabilistic logic program defines a probability distribution over normal logic programs (called **instances** or **possible worlds** or simply **worlds**)
- The distribution is extended to a joint distribution over worlds and interpretations (or queries)
- The probability of a query is obtained from this distribution

PLP under the Distribution Semantics

- A PLP language under the distribution semantics with a general syntax is **Logic Programs with Annotated Disjunctions** (LPADs)
- Heads of clauses are disjunctions in which each atom is annotated with a probability.
- LPAD T with n clauses: $T = \{C_1, \dots, C_n\}$.
- Each clause C_i takes the form:

$$h_{i1} : \pi_{i1}; \dots; h_{iv_i} : \pi_{iv_i} :- b_{i1}, \dots, b_{iu_i}$$

,

- Each grounding $C_i\theta_j$ of a clause C_i corresponds to a random variable X_{ij} with values $\{1, \dots, v_i\}$
- The random variables X_{ij} are independent of each other.

Example

- UW-CSE domain:

advisedby(A, B) : 0.3 : –

student(A), professor(B), project(C, A), project(C, B).

advisedby(A, B) : 0.6 : –

student(A), professor(B), ta(C, A), taughtby(C, B).

Distribution Semantics

- Case of no function symbols: finite Herbrand universe, finite set of groundings of each clause
- **Atomic choice**: selection of the k -th atom for grounding $C_i\theta_j$ of clause C_i
- Represented with the triple (C_i, θ_j, k)
- Example $C_1 = \text{advisedby}(A, B)$:
0.3 :– $\text{student}(A), \text{professor}(B), \text{project}(C, A), \text{project}(C, B),$
 $(C_1, \{A/\text{bob}, B/\text{harry}, C/p1\}, 1)$
- **Composite choice** κ : consistent set of atomic choices
- The probability of composite choice κ is

$$P(\kappa) = \prod_{(C_i, \theta_j, k) \in \kappa} \pi_{ik}$$

Distribution Semantics

- **Selection** σ : a total composite choice (one atomic choice for every grounding of each clause)
- A selection σ identifies a logic program w_σ called **world**
- The probability of w_σ is $P(w_\sigma) = P(\sigma) = \prod_{(C_i, \theta_j, k) \in \sigma} \pi_{ik}$
- Finite set of worlds: $W_T = \{w_1, \dots, w_m\}$
- $P(w)$ distribution over worlds: $\sum_{w \in W_T} P(w) = 1$

Distribution Semantics

- Ground query q
- We consider only *sound* LPADs, where each possible world has a total well-founded model, so $w \models q$ means that the query q is true in the well-founded model of the program w .
- $P(q|w) = 1$ if q is true in w and 0 otherwise
- $P(q) = \sum_w P(q, w) = \sum_w P(q|w)P(w) = \sum_{w \models q} P(w)$

Example

- UW-CSE domain: the objective is to predict the “advised by” relation between students and professors.

$advisedby(A, B) : 0.3 : -$
 $student(A), professor(B), project(C, A), project(C, B).$

$advisedby(A, B) : 0.6 : -$
 $student(A), professor(B), ta(C, A), taughtby(C, B).$

- $q = advisedby(harry, ben)$ where *harry* is a student, *ben* is a professor, they have one joint project and *ben* teaches one course where *harry* is a TA. So

$$P(advisedby(harry, ben)) = 0.3 \cdot 0.6 + 0.3 \cdot 0.4 + 0.7 \cdot 0.6 = 0.72$$

Hierarchical PLP

- We want to compute the probability of atoms for a predicate $r: r(\mathbf{t})$, where \mathbf{t} is a vector of constants.
- $r(\mathbf{t})$ can be an example in a learning problem and r a **target predicate**.
- A specific form of an LPADs defining r in terms of the input predicates.
- The program defined r using a number of input and **hidden predicates** disjoint from input and target predicates.
- Each rule in the program has a single head atom annotated with a probability.
- The program is hierarchically defined so that it can be divided into layers.

Hierarchical PLP

- Each layer contains a set of hidden predicates that are defined in terms of predicates of the layer immediately below or in terms of input predicates.
- Extreme form of program stratification: stronger than **acyclicity** [Apt NGC91] because it is imposed on the predicate dependency graph, and is also stronger than **stratification** [Chandra, Harel JLP85] that allows clauses with positive literals built on predicates in the same layer.
- It prevents inductive definitions and recursion in general, thus making the language not Turing-complete.

Hierarchical PLP

- Generic clause C :

$$C = p(\mathbf{X}) : \pi : - \phi(\mathbf{X}, \mathbf{Y}), b_1(\mathbf{X}, \mathbf{Y}), \dots, b_m(\mathbf{X}, \mathbf{Y})$$

where $\phi(\mathbf{X}, \mathbf{Y})$ is a conjunction of literals for the input predicates using variables \mathbf{X}, \mathbf{Y} .

- $b_i(\mathbf{X}, \mathbf{Y})$ for $i = 1, \dots, m$ is a literal built on a hidden predicate.
- \mathbf{Y} is a possibly empty vector of variables existentially quantified with scope the body.
- Literals for hidden predicates must use the whole set of variables \mathbf{X}, \mathbf{Y} .
- The predicate of each $b_i(\mathbf{X}, \mathbf{Y})$ does not appear elsewhere in the body of C or in the body of any other clause.

Hierarchical PLP

- A generic program defining r is thus:

$$C_1 = r(\mathbf{X}) : \pi_1 \quad :- \quad \phi_1, b_{11}, \dots, b_{1m_1}$$

...

$$C_n = r(\mathbf{X}) : \pi_n \quad :- \quad \phi_n, b_{n1}, \dots, b_{nm_n}$$

$$C_{111} = r_{11}(\mathbf{X}) : \pi_{111} \quad :- \quad \phi_{111}, b_{1111}, \dots, b_{111m_{111}}$$

...

$$C_{11n_{11}} = r_{11}(\mathbf{X}) : \pi_{11n_{11}} \quad :- \quad \phi_{11n_{11}}, b_{11n_{11}1}, \dots, b_{11n_{11}m_{11n_{11}}}$$

...

$$C_{n11} = r_{n1}(\mathbf{X}) : \pi_{n11} \quad :- \quad \phi_{n11}, b_{n111}, \dots, b_{n11m_{n11}}$$

...

$$C_{n1n_{n1}} = r_{n1}(\mathbf{X}) : \pi_{n1n_{n1}} \quad :- \quad \phi_{n1n_{n1}}, b_{n1n_{n1}1}, \dots, b_{n1n_{n1}m_{n1n_{n1}}}$$

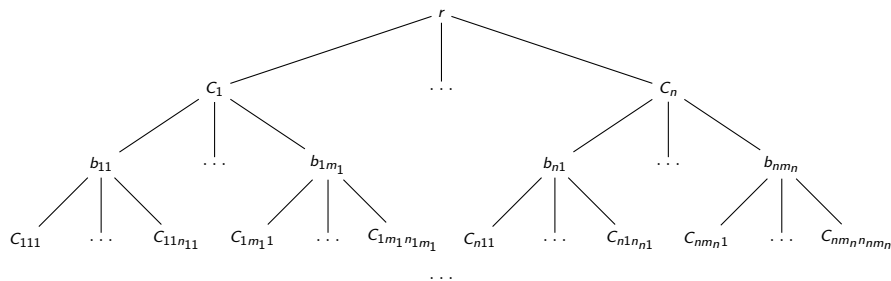
...



Example


- $$\begin{aligned} C_1 &= \textit{advisedby}(A, B) : 0.3 :- \\ &\quad \textit{student}(A), \textit{professor}(B), \textit{project}(C, A), \textit{project}(C, B), \\ &\quad r_{11}(A, B, C). \\ C_2 &= \textit{advisedby}(A, B) : 0.6 :- \\ &\quad \textit{student}(A), \textit{professor}(B), \textit{ta}(C, A), \textit{taughtby}(C, B). \\ C_{111} &= r_{11}(A, B, C) : 0.2 :- \\ &\quad \textit{publication}(D, A, C), \textit{publication}(D, B, C). \end{aligned}$$

Program Tree



Example

- $C_1 = \text{advisedby}(A, B) : 0.3 :-$
 $\text{student}(A), \text{professor}(B), \text{project}(C, A), \text{project}(C, B),$
 $r_{11}(A, B, C).$
- $C_2 = \text{advisedby}(A, B) : 0.6 :-$
 $\text{student}(A), \text{professor}(B), \text{ta}(C, A), \text{taughtby}(C, B).$
- $C_{111} = r_{11}(A, B, C) : 0.2 :-$
 $\text{publication}(D, A, C), \text{publication}(D, B, C).$



Hierarchical PLP

- Writing programs in hierarchical PLP may be unintuitive for humans because of the need of satisfying the constraints and because the hidden predicates may not have a clear meaning.
- The structure of the program should be learned by means of a specialized algorithm
- Hidden predicates generated by a form of predicate invention.

Inference

- Generate the grounding.
- Each ground probabilistic clause is associated with a random variable whose probability of being true is given by the parameter of the clause and that is independent of all the other clause random variables.
- Ground clause $C_{\mathbf{p}i} = a_{\mathbf{p}} : \pi_{\mathbf{p}i} :- b_{\mathbf{p}i1}, \dots, b_{\mathbf{p}im_{\mathbf{p}}}$. where \mathbf{p} is a path in the program tree
- $P(b_{\mathbf{p}i1}, \dots, b_{\mathbf{p}im_{\mathbf{p}}}) = \prod_{i=k}^{m_{\mathbf{p}}} P(b_{\mathbf{p}ik})$ and $P(b_{\mathbf{p}ik}) = 1 - P(a_{\mathbf{p}ik})$ if $b_{\mathbf{p}ik} = \neg a_{\mathbf{p}ik}$.
- If a is a literal for an input predicate, then $P(a) = 1$ if a belongs to the example interpretation and $P(a) = 0$ otherwise.

Inference


- Hidden predicates: to compute $P(a_p)$ we need to take into account the contribution of every ground clause for the predicate of a_p .
- Suppose these clauses are $\{C_{p1}, \dots, C_{po_p}\}$.
- If we have two clauses,
$$P(a_{pi}) = 1 - (1 - \pi_{p1} \cdot P(\text{body}(C_{p1}))) \cdot (1 - \pi_{p2} \cdot P(\text{body}(C_{p2})))$$
- $p \oplus q \triangleq 1 - (1 - p) \cdot (1 - q)$.
- This operator is commutative and associative:

$$\bigoplus_i p_i = 1 - \prod_i (1 - p_i)$$

- The operators \times and \oplus are respectively the t-norm and t-conorm of the product fuzzy logic [Hajek 98]: **product t-norm** and **probabilistic sum**.

Inference

- If the probabilistic program is ground, the probability of the example atom can be computed with the arithmetic circuit:





- The arithmetic circuit can be interpreted as a deep neural network where nodes have the activation functions \times and \oplus

Example

- $G_1 = \text{advisedby}(\text{harry}, \text{ben}) : 0.3 : -$
 $\text{student}(\text{harry}), \text{professor}(\text{ben}), \text{project}(\text{pr1}, \text{harry}),$
 $\text{project}(\text{pr1}, \text{ben}), r_{11}(\text{harry}, \text{ben}, \text{pr1}).$
- $G_2 = \text{advisedby}(\text{harry}, \text{ben}) : 0.3 : -$
 $\text{student}(\text{harry}), \text{professor}(\text{ben}), \text{project}(\text{pr2}, \text{harry}),$
 $\text{project}(\text{pr2}, \text{ben}), r_{11}(\text{harry}, \text{ben}, \text{pr2}).$
- $G_3 = \text{advisedby}(\text{harry}, \text{ben}) : 0.6 : -$
 $\text{student}(\text{harry}), \text{professor}(\text{ben}), \text{ta}(\text{c1}, \text{harry}), \text{taughtby}(\text{c1}, \text{ben}).$
- $G_4 = \text{advisedby}(\text{harry}, \text{ben}) : 0.6 : -$
 $\text{student}(\text{harry}), \text{professor}(\text{ben}), \text{ta}(\text{c2}, \text{harry}), \text{taughtby}(\text{c2}, \text{ben}).$
- $G_{111} = r_{11}(\text{harry}, \text{ben}, \text{pr1}) : 0.2 : -$
 $\text{publication}(\text{p1}, \text{harry}, \text{pr1}), \text{publication}(\text{p1}, \text{ben}, \text{pr1}).$
- $G_{112} = r_{11}(\text{harry}, \text{ben}, \text{pr1}) : 0.2 : -$
 $\text{publication}(\text{p2}, \text{harry}, \text{pr1}), \text{publication}(\text{p2}, \text{ben}, \text{pr1}).$
- $G_{211} = r_{11}(\text{harry}, \text{ben}, \text{pr2}) : 0.2 : -$
 $\text{publication}(\text{p3}, \text{harry}, \text{pr2}), \text{publication}(\text{p3}, \text{ben}, \text{pr2}).$
- $G_{212} = r_{11}(\text{harry}, \text{ben}, \text{pr2}) : 0.2 : -$
 $\text{publication}(\text{p4}, \text{harry}, \text{pr2}), \text{publication}(\text{p4}, \text{ben}, \text{pr2}).$

Example



Building the Network

- The network can be built by performing inference using tabling and answer subsumption
- PITA(IND,IND) [Riguzzi CJ14] is a program transformation that adds an extra argument to each subgoal of the program and of the query to store the probability of answers to the subgoal
- When a subgoal returns, the extra argument will be instantiated to the probability of the ground atom that corresponds to the subgoal without the extra argument.
- In programs of hierarchical PLP, when a subgoal returns the original arguments are guaranteed to be instantiated.
- PITA(IND,IND) adds literals to bodies that combine the extra arguments of the subgoals

Building the Network

- The contributions of multiple groundings of multiple clauses are combined by means of tabling with answer subsumption.
- Tabling: keep a store of the subgoals encountered in a derivation together with answers to these subgoals.
- If one of the subgoals is encountered again, its answers are retrieved from the store rather than recomputing them.
- Tabling reduces computation time and ensures termination for a large class of programs [Swift TPLP12].
- Answer subsumption [Swift TPLP12] is a tabling feature that, when a new answer for a tabled subgoal is found, combines old answers with the new one.
- In PITA(IND, IND) the combination operator is probabilistic sum.

Parameter Learning

- Parameter learning by EM or backpropagation.
- Inference has to be performed repeatedly on the same program with different values of the parameters.
- PITA(IND,IND) can build a representation of the arithmetic circuit, instead of just computing the probability.
- Extra argument used to store a term representing the circuit
- Implementing EM would adapt the algorithm of [Bellodi and Riguzzi IDA13] for hierarchical PLP.

Related Work

- [Sourek et al NIPS15]: build deep neural networks using a template expressed as a set of weighted rules.
- Nodes for ground atoms and ground rules
- Values of ground rule nodes aggregated to compute the value of atom nodes.
- Aggregation in two steps, first the contributions of different groundings of the same rule sharing the same head and then the contributions of groundings for different rules.
- Proposal parametric in the activation functions of ground rule nodes.
- Example: two families of activation functions that are inspired by Lukasiewicz fuzzy logic.
- We build a neural network whose output is the probability of the example according to the distribution semantics.

Related Work

- **Edward** [Tran et al. ICLR17]: Turing-complete probabilistic programming language
- Programs in Edward define computational graphs and inference is performed by stochastic graph optimization using TensorFlow.
- Hierarchical PLP is not Turing-complete as Edward but ensures fast inference by circuit evaluation.
- Being based on logic it handles well domains with multiple entities connected by relationships.
- Similarly to Edward, hierarchical PLP can be compiled to TensorFlow

Related Work

- Probabilistic Soft Logic (PSL) [Bach et al. arXiv15]: Markov Logic with atom random variables taking continuous values in $[0, 1]$ and logic formulas interpreted using Lukasiewicz fuzzy logic.
- PSL defines a joint probability distribution over fuzzy variables, while the random variables in hierarchical PLP are still Boolean and the fuzzy values are the probabilities that are combined with the product fuzzy logic.
- The main inference problem in PSL is MAP rather than MARG as in hierarchical PLP.

Related Work

- Sum-product networks [Poon, Domingos UAI11]: hierarchical PLP circuits can be seen as sum-product networks where children of sum nodes are not mutually exclusive but independent and each product node has a leaf child that is associated to a hidden random variable.
- Sum-product networks represent a distribution over input data while programs in hierarchical PLP describe only a distribution over the truth values of the query.
- Inference in hierarchical PLP is in a way “lifted”: the probability of the ground atoms can be computed knowing only the sizes of the populations of individuals that can instantiate the existentially quantified variables

Conclusions and Future Work

- Conclusions
 - hierarchical PLP: a restriction of the language of LPADs that allows to perform inference quickly using a simple and cheap dynamic programming algorithm such as PITA(IND,IND).
 - Programs can be seen as arithmetic circuits/neural networks
 - Parameters can be trained by gradient descent or EM
- Future work
 - Develop backpropagation and EM formulas
 - Perform also structure learning



**THANKS FOR
LISTENING
AND
ANY
QUESTIONS ?**