

Probabilistic Logic Learning

Luc De Raedt

Institut für Informatik, Albert-Ludwigs-University,
Georges-Koehler-Allee, Building 079, D-79110
Freiburg, Germany

deraedt@informatik.uni-freiburg.de

Kristian Kersting

Institut für Informatik, Albert-Ludwigs-University,
Georges-Koehler-Allee, Building 079, D-79110
Freiburg, Germany

kersting@informatik.uni-freiburg.de

ABSTRACT

The past few years have witnessed an significant interest in *probabilistic logic learning*, i.e. in research lying at the intersection of probabilistic reasoning, logical representations, and machine learning. A rich variety of different formalisms and learning techniques have been developed. This paper provides an introductory survey and overview of the state-of-the-art in probabilistic logic learning through the identification of a number of important probabilistic, logical and learning concepts.

Keywords

data mining, machine learning, inductive logic programming, multi-relational data mining, uncertainty, probabilistic reasoning

1. INTRODUCTION

One of the central open questions in data mining and artificial intelligence, concerns *probabilistic logic learning*, i.e. the integration of relational or logical representations, probabilistic reasoning mechanisms with machine learning and data mining principles. In the past few years, this question has received a lot of attention and various different approaches have been developed, cf. [82; 41; 88; 66; 73; 45; 60; 61; 57; 81; 52; 77; 2; 55; 87]. This paper provides an introductory survey and overview of those developments that lie at the intersection of logical (or relational) representations, probabilistic reasoning and learning, cf. Figure 1. While doing so, we identify some of the key concepts and techniques underlying probabilistic logic learning. This, in turn, we hope, should allow the reader to appreciate the differences and commonalities between the various approaches and at the same time provide insight into some of the remaining challenges in probabilistic logic learning.

Before starting our survey, let us specify what we mean by *probabilistic logic learning*. The term *probabilistic* in our context refers to the use of probabilistic representations and reasoning mechanisms grounded in probability theory, such as *Bayesian networks* [78; 10; 47], *hidden Markov models* [85; 84; 22] and *stochastic grammars* [22; 64]. Such representations have been successfully used across a wide range of applications and have resulted in a number of robust models for reasoning about uncertainty. Application areas include genetics, computer vision, speech recognition and understand-




Figure 1: Probabilistic Logic Learning as the intersection of *Probability*, *Logic*, and *Learning*.

ing, diagnostic and troubleshooting, information retrieval, software debugging, data mining and user modelling.

The term *logic* in the present overview refers to first order logical and relational representations such as those studied within the field of computational logic. The primary advantage of using *first order* logic or relational representations is that it allows one to elegantly represent complex situations involving a variety of objects as well as relations between the objects. Consider e.g. the problem of building a logical troubleshooting system of a set of computers. Computers are complex objects which are themselves composed of several complex objects. Furthermore, assume that even though the overall structure of the computers can be the same, their individual components can be highly different. Such situations can be elegantly modelled using relational or first order logic representations. E.g., we could use the relation $\text{ram}(\mathbf{r}_1, \mathbf{c}_1)$ to specify that computer \mathbf{c}_1 has a RAM component of type \mathbf{r}_1 , and $\text{cpu}(\mathbf{p}_2, \mathbf{c}_1)$ and $\text{cpu}(\mathbf{p}_3, \mathbf{c}_1)$ to denote that \mathbf{c}_1 is a multi-processor machine with two cpu's \mathbf{p}_2 and \mathbf{p}_3 . In addition to these base (or extensional) relations, one can also add virtual (or intensional) relations to describe generic knowledge about the domain. E.g., consider the rule

$$\text{multiprocessor}(\mathbf{C}) \text{ :- } \text{cpu}(\mathbf{P1}, \mathbf{C}), \text{cpu}(\mathbf{P2}, \mathbf{C}), \mathbf{P1} \neq \mathbf{P2}$$

which defines the concept of a multi-processor machine (as a machine that possesses two different cpu's). Representing the same information employing a propositional logic would require one to specify for each machine a separate model.

The term *learning* in the context of probabilistic logic refers to deriving the different aspects of the probabilistic logic on the basis of data. Typically, one distinguishes various learning algorithms on the basis of the given data (fully or partially observable variables) or on the aspect being learned (the parameters of the probabilistic representation or its logical structure). The motivation for learning is that it is of-

ten easier to obtain data for a given application domain and learn the model than to build the model using traditional knowledge engineering techniques.

So, probabilistic logic learning aims at combining its three underlying constituents: learning and probabilistic reasoning within first order logic representations. The recent interest in probabilistic logic learning may be explained by the growing body of work that has addressed pairwise intersections of these domains:

- Indeed, various techniques for *probabilistic learning* such as gradient-based methods, the family of EM algorithms or Markov Chain Monte Carlo methods have been developed and exhaustively investigated in different communities, such as in the Uncertainty in AI community for Bayesian networks and in the Computational Linguistic community for Hidden Markov Models. These techniques are not only theoretically sound, they have also resulted in entirely new technologies for, and revolutionary novel products in computer vision, speech recognition, medical diagnostics, troubleshooting systems, etc. Overviews of and introductions to probabilistic learning can be found in [20; 84; 44; 8; 48; 22; 64].
- Inductive Logic Programming and multi-relational data mining studied *logic learning*, i.e. learning and data mining within first order logical or relational representations. Inductive Logic Programming has significantly broadened the application domain of data mining especially in bio- and chemo-informatics (cf. the other papers in this volume) and now represent some of the best-known examples of scientific discovery by AI systems in the literature. Overviews of inductive logic learning and multi-relational data mining can be found in this volume and in [69; 23].
- Early work on incorporating probabilities into logic programs was done by Clark and McCabe [9] and by Shapiro [97]. Subsequently, researchers such as Nilsson [75], Halpern [43], Ng and Subrahmanian [71], and Poole [82] have studied *probabilistic logics* from a knowledge representational perspective. The aim of this research is more a probabilistic characterization of logic than suitable representations for learning.

Finally, in the past few years there have been some important attempts that address the full problem of probabilistic logic learning [59; 30; 67; 31; 14; 53; 51; 92; 68].

The goal of this paper is to provide an introduction and overview to these works. In doing so, we restrict our attention to those developments that address the intersection of probabilistic logic learning, cf. Figure 1. Consequently, we will not discuss interesting work lying on boundaries such as Flach and Lachiche’s work on Naïve Bayes for structured terms [27; 62] and Craven and Slattery’s work on combining Naïve Bayes with a relational rule learner [11]. Giving an overview of the pairwise intersections (let alone the underlying domains) would lead us too far (as it would require a book instead of a paper) and is therefore beyond the scope of this paper. Because probabilistic logic learning involves three domains, one can approach this topic from various directions. In this paper, motivated by the topic of this special issue, we start from a multi-relational and inductive logic programming perspective. The terminology and notation employed throughout this paper is based on computational logic rather than on relational databases because

```
alarm      :- burglary, earthquake.
johncalls :- alarm.
marycalls  :- alarm.
```

Figure 2: The **alarm** program.

the relational model is not expressive enough to model the logical component of many of the probabilistic logical representations (e.g., [88; 66; 90; 52]). Nevertheless, we have attempted to be self-contained and to reduce the required logical machinery and background as much as possible.

The paper is organized as follows: Sections 2 and 3 introduce the key concepts underlying computational logic and probabilistic reasoning formalisms, respectively; Section 4 then studies first-order probabilistic logics, i.e., how probabilistic reasoning and computational logic can be combined; afterwards, Section 5 then surveys various approaches to learning within those probabilistic logics, and, finally, in Section 6, we conclude.

2. INTRODUCTION TO LOGIC

As logic programs will be used throughout this paper as the underlying knowledge representation framework, we now briefly introduce their key concepts. Afterwards, we show how state-of-the-art probabilistic representations employ these concepts.

2.1 Logic

The choice of logic programs as the underlying representation is motivated by its expressive power and generality. Indeed, logic programs allow us to represent relational databases, grammars and also programs. This will turn out to be useful when discussing commonalities and differences in the probabilistic relational or logical representations later on. In addition, logic programs are well-understood, have a clear semantics and are wide-spread.

Three example logic programs are given in Figures 2, 3 and 4.b. The first one specifies some regularities in the **alarm** program (inspired on Judea Pearl’s famous example for Bayesian networks). The **genetic** program is a simple model of Mendel’s law applied to the color of plants. Finally, the **grammar** example encodes a context-free grammar. The corresponding type of logic program is known under the term “definite clause grammar”.

Let us now introduce some terminology. The *predicates* in the above examples include **alarm/0**, **burglary/0**, **earthquake/0**, ..., **pc/2**, **mc/2**, **pt/2**, **mother/2**, **father/2**, **s/2**, **vp/2**, ... Their *arity* (i.e., the number of arguments) has been listed explicitly. Predicates with arity 0, such as **alarm/0** are sometimes called *propositions*. Furthermore, **1**, **1m**, **1mf**, ..., **g**, **green**, ..., **dog**, **barks**, ... are *constants* and **P**, **M**, **X**, **C**, **S**, **T**, ... are *variables*. Constants and variables are also *terms*. In addition, there exist structured terms such as **[barks]**, shorthand for **[] (barks, nil)**, which contains the *list functor* **[]/2**, and the terms **barks** and **nil** (the empty list). Constants are often considered as functors of arity 0. *Atoms* are predicate symbols followed by the necessary number of terms, e.g., **v([barks|T], T)** and **pt(P, white)**. We are now able to define the key concept of a *definite clause*. Definite clauses are formulas of the form **A :- B₁, ..., B_m** where **A** and the **B_i** are atoms and where all variables are understood to be universally quantified. E.g.,

```

mc(P, X)      :- mother(P, M), mc(M, X), pc(M, Y).
mc(P, X)      :- mother(P, M), mc(M, Y), pc(M, X).
pc(P, X)      :- father(P, F), mc(F, X), pc(F, Y).
pc(P, X)      :- father(P, F), mc(F, Y), pc(F, X).
mc(1mf, r).   mother(1, 1m).
pc(1mf, g).   father(1, 1f).
mc(1mm, g).   mother(1m, 1mm).
pc(1mm, r).   father(1m, 1mf).
mc(1f, g).
pc(1f, g).
pt(P, green) :- mc(P, g), pc(P, g).
pt(P, red)   :- mc(P, r).
pt(P, red)   :- pc(P, r).

```

Figure 3: The **genetic** program.

S	→ NP, VP.	s(S, T) :- np(S, S1), vp(S1, T).
VP	→ V.	vp(S, T) :- v(S, T).
NP	→ Det, Noun.	np(S, T) :- det(S, S1), noun(S1, T).
Det	→ [the].	det([the T], T).
Noun	→ [dog].	noun([dog T], T).
V	→ [barks].	v([barks T], T).
(a)		(b)

Figure 4: Context-free grammars: (a) A context-free grammar and (b) the corresponding **grammar** program.

the definite clause $\text{pt}(\text{P}, \text{green}) :- \text{mc}(\text{P}, \text{g}), \text{pc}(\text{P}, \text{g})$ can be read as “the phenotype of P is green if the m-chromosome of P contains the gene g and the p-chromosome of P contains the gene g”. Because we will restrict attention to definite clauses only in the present paper, we will omit the term “definite” as long as no ambiguities arise. We call $\text{pt}(\text{P}, \text{green})$ the *head* of the clause, and $\text{mc}(\text{P}, \text{g}), \text{pc}(\text{P}, \text{g})$ the *body*. Clauses with an empty body, such as $v([\text{barks}|T], T)$ are called *facts*. A clause program (or *logic program* for short) consists of a set of clauses. For a more thorough introduction to logic programming we refer to [63; 26].

Within logic, two types of semantics can be distinguished: a *model-theoretic* and a *proof-theoretic* one. Because this distinction will also be useful from a probabilistic point of view, let us briefly introduce these two concepts.

2.2 A Model Theoretic View

Consider the **alarm** program in Figure 2. In this example, there are five propositions that occur, i.e. $\{\text{alarm}, \text{earthquake}, \text{marycalls}, \text{johncalls}, \text{burglary}\}$. This is the so-called *Herbrand base* **HB**. It consists of all facts that one can construct with the predicate, constant and function symbols in the program. The Herbrand base – in a sense – specifies the set of all possible worlds described by the program. Indeed, for the **alarm** program, there are $2^5 = 32$ possible assignments of truth-values to the Herbrand base. These assignments are sometimes called (*Herbrand*) *interpretations* and they each describe a world.

From a model-theoretic point of view, a logic program restricts the set of possible worlds. Indeed, if we interpret

the clause $\text{alarm} :- \text{burglary}, \text{earthquake}$ as an implication, then the interpretations in which both **burglary** and **earthquake** are true, but **alarm** is not, do not satisfy this clause. The idea is that if we accept the clause as an axiom, then such worlds are impossible. More formally, we have that a (Herbrand) interpretation I satisfies a clause c if and only if for all substitutions θ such that $\text{body}(c)\theta \subseteq I$ also $\text{head}(c)\theta \in I$. The interpretation I is called a model of c if it satisfies c . A *substitution* $\theta = \{V_1/t_1, \dots, V_n/t_n\}$, e.g. $\theta_1 = \{\text{P}/1\text{mm}\}$, is an assignment of terms t_i to variables V_i . Applying a substitution θ to a term, atom or clause e yields the instantiated term, atom, or clause $e\theta$ where all occurrences of the variables V_i are simultaneously replaced by the term t_i . Consider e.g. the clause $c \text{pt}(\text{P}, \text{green}) :- \text{mc}(\text{P}, \text{g}), \text{pc}(\text{P}, \text{g})$ and the substitution θ_1 . Applying θ_1 to c yields the instantiated clause $c\theta_1 \text{pt}(1\text{mm}, \text{green}) :- \text{mc}(1\text{mm}, \text{g}), \text{pc}(1\text{mm}, \text{g})$. Furthermore, the reader may want to verify that the interpretation $\{\text{mc}(1\text{m}, \text{g}), \text{pc}(1\text{m}, \text{g})\}$ is not a model for the clause c . Finally, let us remark that an interpretation satisfies a set of clauses if it satisfies all the clauses in the set.

The most important point about model theory – for our purposes – is that it specifies which worlds (i.e., interpretations) are possible with respect to (i.e., satisfy) a given logical theory.

2.3 A Proof Theoretic View

Another way of looking at logic programs comes from proof theory. From this perspective, a logic program can be used to prove that certain atoms, goals (see below) or clauses are logically entailed by the program. Consider, e.g., the **genetic** program in Figure 3. In this program, the fact $\text{pt}(\text{P}, \text{C})$ will be provable for any P and C whenever C is one of the possible colors for the plant P according to the axioms encoded in the program.

Proofs are typically constructed using SLD-resolution procedure, which we now briefly introduce. More formally, given a goal $:- G_1, \dots, G_n$ and a clause $G :- L_1, \dots, L_m$ such that $G_1\theta = G\theta$, applying resolution results in the goal $:- L_1\theta, \dots, L_m\theta, G_2\theta, \dots, G_n\theta$. A (successful) resolution refutation of a goal $:- G$ is then a sequence of resolution steps that leads to the empty goal $:-$. A failure does not end in the empty goal. For instance, in the genetic example, one can prove that $\text{pt}(1\text{f}, \text{green})$ is true, because of the following resolution derivation

```

:- pt(1f, green).
:- mc(1f, g), pc(1f, g).
:- pc(1f, g).
:-

```

Resolution is employed by many theorem provers (such as Prolog). Indeed, when given the goal $:- \text{pt}(1\text{f}, \text{green})$, Prolog would compute the above resolution derivation and answer *yes*. Furthermore, when given the goal $:- \text{pt}(1\text{f}, \text{X})$, Prolog would answer with the substitutions that make the goal true, i.e., $\{\text{X}/\text{green}\}$. The set of possible resolution steps for a given goal is captured in an SLD-tree. Two examples, one in the genetic domain, and one in the grammar domain are shown in Figure 5. At this point it is important to see the connection between an SLD-derivation for the grammar domain and a traditional derivation using the context free grammar. Indeed, there is a one to one mapping between these two concepts for our grammar.

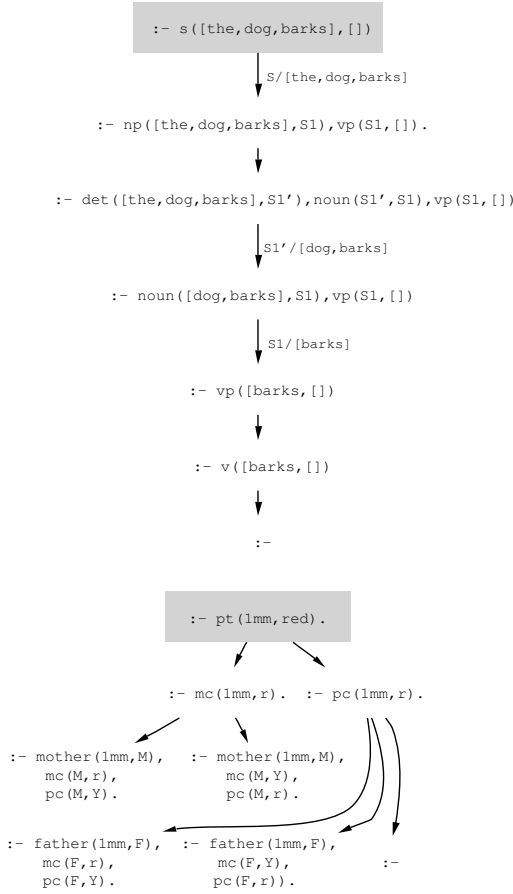


Figure 5: Examples of SLD trees. Root nodes are shaded. The upper tree is the SLD tree of querying the **grammar** program with `:- s([the,dog,barks],[])`. The lower tree is the SLD tree of querying the **genetic** program with `:- pt(1mm,red)`. In both cases, there is only one refutation, i.e. successful derivation.

There are thus two ways of viewing a logic program. The first view, illustrated by the **alarm** program, is the *model-theoretic* view. It determines what the possible worlds (i.e. interpretations) are that satisfy the program. The second view, illustrated by the **genetic** program, is the *proof-theoretic* view. It determines what the possible proofs are with respect to a given goal. Of course, these two views are intimately connected. Indeed, for clause programs, we have that a (ground) fact has a resolution refutation if and only if it belongs to all (Herbrand) models of the program.

3. PROBABILISTIC LOGICS

The model theoretic and proof theoretic views of logic are also useful from a probabilistic perspective. Indeed, many of the present probabilistic representations can be regarded as defining probabilities on possible worlds (e.g., Bayesian networks) or on proofs (e.g., stochastic context free grammars). As these probabilistic representations form the basis for the probabilistic logics, we will now provide a short overview to these. We will use \mathbf{P} to denote a probability distribution, e.g. $\mathbf{P}(X)$, and P to denote a probability value, e.g. $P(X = x)$.




Figure 6: The graphical structure of the alarm Bayesian network.

3.1 Probabilities on Possible Worlds

The most popular formalism for defining probabilities on possible worlds is that of Bayesian networks. As an example of a Bayesian network, consider Judea Pearl's famous alarm network graphically illustrated in Figure 6. Formally speaking, a *Bayesian network* [78; 10; 47] is an augmented, directed acyclic graph, where each node corresponds to a random variable X_i (e.g., $\{\text{alarm}, \text{earthquake}, \text{marycalls}, \text{johncalls}, \text{burglary}\}$) and each edge indicates a direct influence among the random variables. We will denote the set of parents of a node X_i by $\mathbf{Pa}(X_i)$, e.g., $\mathbf{Pa}(\text{alarm}) = \{\text{earthquake}, \text{burglary}\}$. A Bayesian network specifies a joint probability distribution $\mathbf{P}(X_1, \dots, X_n)$ over a fixed, finite set $\{X_1, \dots, X_n\}$ of random variables. As we will – for simplicity – assume that the random variables are all boolean, i.e., they can have the domain $\{\text{true}, \text{false}\}$ ¹, this actually amounts to specifying a probability distribution on the set of all possible interpretations. Indeed, in our alarm example, the Bayesian network defines a probability distribution over truth-assignments to $\{\text{alarm}, \text{earthquake}, \text{marycalls}, \text{johncalls}, \text{burglary}\}$.

A Bayesian network stipulates the following conditional independency assumption:

Each node X_i in the graph is conditionally independent of any subset \mathbf{A} of nodes that are not descendants of X_i given a joint state of $\mathbf{Pa}(X_i)$, i.e. $\mathbf{P}(X_i \mid \mathbf{A}, \mathbf{Pa}(X_i)) = \mathbf{P}(X_i \mid \mathbf{Pa}(X_i))$.

For example, **marycalls** is conditionally independent of **earthquake** and of **burglary** given a joint state of **alarm**. Because of the conditional independence assumption, we can write down the joint probability density as

$$\mathbf{P}(X_1, \dots, X_n) = \prod_{i=1}^n \mathbf{P}(X_i \mid \mathbf{Pa}(X_i))$$

by applying the independency assumption to the chain rule expression of the joint probability distribution. Therefore, we associate with each node X_i of the graph the conditional probability distribution $\mathbf{P}(X_i \mid \mathbf{Pa}(X_i))$, denoted as $\text{cpd}(X_i)$. The conditional probability distributions in our alarm examples would include $\mathbf{P}(\text{alarm} \mid \text{earthquake}, \text{burglary})$, $\mathbf{P}(\text{earthquake})$, etc. They are illustrated in Table 1.

So, Bayesian networks define a probability distribution on the possible worlds, i.e., the interpretations. Furthermore, the resulting probabilities induce a probability distribution at the level of propositions and propositional formulae. Indeed, reconsider the **alarm** example, and assume that we

¹In general, the domains of a random variable can be continuous or discrete as well, cf. [25].

$P(\text{burglary})$		$P(\text{earthquake})$	
(0.001, 0.999)		(0.002, 0.998)	

burglary	earthquake	$P(\text{alarm})$
true	true	(0.95, 0.05)
true	false	(0.94, 0.06)
false	true	(0.29, 0.71)
false	false	(0.001, 0.999)

alarm	$P(\text{johncalls})$	alarm	$P(\text{marycalls})$
true	(0.90, 0.10)	true	(0.70, 0.30)
false	(0.05, 0.95)	false	(0.01, 0.99)

Table 1: The conditional probability distributions associated with the nodes in the alarm network, cf. Figure 6. The distributions are specified over $\{true, false\}$. When viewing the child-parents pairs as clauses, the clauses in the **alarm** program, cf. Figure 2, specify some regularities of the alarm network. The corresponding conditional probability distributions can be associated with the clauses.

are interested in the probability of $P(\text{alarm} = true)$. This probability can be obtained by marginalizing over the other variables, i.e.,

$$P(a = t) = \sum_{x_i \in \{t, f\}} P(a = t, b = x_1, e = x_2, m = x_3, j = x_4)$$

where we have abbreviated the variables and states appropriately. The probability of a propositional formula can be obtained similarly, it is simply the sum of the probabilities of the interpretations that are a model for the formula.

The key limitation of Bayesian networks is that they lack the notion of an object. The possible worlds considered are *propositional*. Using Bayesian networks, it is impossible to model relations between objects, i.e., relational or logical worlds. This problem is similar to that with traditional propositional mining and learning systems.

3.2 Probabilities on Proofs

Because of the closeness of proofs of logical formulae and derivations in grammar, we will start by discussing stochastic grammars, see e.g. [1; 64]. Consider therefore the stochastic context free grammar in Figure 7 (adapted from [1]), where we have omitted many of the terminal symbols for brevity.

This context-free grammar is *stochastic* in that every grammar rule has a probability associated to it. Moreover, for every non-terminal symbol T the sum of the probabilities of the rules with T on the left hand side equals 1.0, e.g., for NP: $0.2 + 0.8 = 1.0$. Stochastic context-free grammars define probabilities on derivations (or, when mapped onto logic programs, on proofs). Indeed, one of the proofs (with associated probabilities) of the sentence "*Rice flies like sand*" goes as follows:

$\vdash S.$	1.0
$\vdash NP VP.$	1.0
$\vdash N N VP.$	0.2
$\vdash \text{Rice } N VP.$	$0.2 \cdot 0.01$
$\vdash \text{Rice flies } VP.$	$0.2 \cdot 0.01 \cdot 0.01$
$\vdash \text{Rice flies } V NP.$	$0.2 \cdot 0.01 \cdot 0.01 \cdot 0.6$

(1.0)	$S \rightarrow NP, VP.$
(0.2)	$NP \rightarrow N, N.$
(0.8)	$NP \rightarrow N.$
(0.6)	$VP \rightarrow V, NP.$
(0.4)	$VP \rightarrow V, PP.$
(1.0)	$PP \rightarrow P, NP.$
(0.01)	$N \rightarrow [\text{Rice}].$
(0.01)	$N \rightarrow [\text{flies}].$
...	
(0.02)	$V \rightarrow [\text{like}].$
(0.02)	$V \rightarrow [\text{flies}].$
...	
(0.05)	$P \rightarrow [\text{like}].$
...	

Figure 7: An example of a stochastic context-free grammar. The numbers associated with the grammar rules denote probability values. Note that $[\text{Rice}]$ does not denote a variable but the string *Rice*.

$\vdash \text{Rice flies like } NP.$	$0.2 \cdot 0.01 \cdot 0.01 \cdot 0.6 \cdot 0.02$
$\vdash \text{Rice flies like } N.$	$0.2 \cdot 0.01 \cdot 0.01 \cdot 0.6 \cdot 0.02 \cdot 0.8$
$\vdash \text{Rice flies like sand.}$	$0.2 \cdot 0.01 \cdot 0.01 \cdot 0.6 \cdot 0.02 \cdot 0.8 \cdot 0.1$

Each time a rule is applied in a proof, the probability of the rule is multiplied with the overall probability. It is easy to verify that this induces a probability distribution over all possible proofs or successful derivations for the same non-terminal symbols. Failures (if any) are only due to non-matching terminals which can be considered to have 0.0 as probability label. Furthermore, the probabilities over the proofs induce a probability over the accepted sentences for each of the non-terminals. Indeed, consider again the sentence "*Rice flies like sand*" for which there are two proofs, one with probability p_1 (as illustrated above) and another one with probability p_2 . The sum of these probabilities specifies the probability that a random sentence generated by the non-terminal symbol S is indeed "*Rice flies like sand*". These distributions are useful for various purposes. For instance, in natural language processing, one may be interested in either the most likely parse tree (or derivation) when dealing with ambiguity, or the total probability that a particular sentence is derived. Also, in bioinformatics, there are numerous applications of stochastic grammars, most notably of (hidden) Markov models [84], which are essentially stochastic regular grammars. These are e.g. used to determine how likely it is that a given protein belongs to some fold [22].

From a computational and logical perspective, the key limitation of stochastic context free grammars concerns their expressive power. Indeed, context free grammars (and therefore stochastic context free grammars) are not Turing-equivalent, i.e. they cannot be used as a programming language. Thus it is useful to lift the underlying grammar representation to that of logic programs, a Turing equivalent language, cf. [66].

4. FIRST ORDER PROBABILISTIC LOGICS

By now, everything is in place to introduce the key representational frameworks that combine probabilistic reasoning with logical or relational representations. Furthermore, in the spirit of the above presentation, we can distinguish them according to whether they define probabilities on interpretations or on proofs.

4.1 Probabilistic Logical Models

The first class of representations extends Bayesian networks with abilities to define probabilities on first order logical or relational interpretations, cf. [82; 41; 73; 45; 57; 81; 31; 52]. Important predecessors of these logical Bayesian networks discussed below, include the work by Breese, Charniak, Bacchus, Goldman, Poole and Wellmann [40; 6; 4; 7; 39]. These predecessors were largely based on the idea of *knowledge based model construction* (KBMC). In knowledge based model construction, a knowledge base (sometimes in the form of an annotated logic program) is used to describe sets of probabilistic models. A query then results in constructing a particular model that can be used to answer the query. Many of these approaches can represent Bayesian networks. However, they are not *direct* upgrades of Bayesian networks in the same way that the work by e.g. Haddawy [41] is. This may also explain why – to the authors’ knowledge – Bayesian network learning techniques have not been applied to these representations.

A central contribution in combining clausal logic with Bayesian networks is due to Peter Haddawy [41]. It is based on the observation that the structure of a Bayesian network can essentially be represented as a set of (propositional) clauses, where there will be one clause for each node in the graph, cf. the **alarm** program and Bayesian network. This observation allows one to upgrade Bayesian networks toward first order logic by upgrading the corresponding propositional clauses towards definite clauses. These clauses will specify the so-called *intensional* part I , i.e. the overall regularities. In addition, we will need an *extensional* part E , which specifies the specific objects and situations under consideration. The distinction between *extensional* and *intensional* is akin to that made in database theory and computational logic. E.g., in the **genetic** program, the extension corresponds to the facts and the intension to the non-fact clauses. Given an extensional part E and an intensional part I , the central ideas are then

- to view the atoms in the Herbrand interpretation as random variables²; they will constitute the nodes in the induced Bayesian network,
- to view the *instantiated* clauses $H :- B_1, \dots, B_m$ as defining the probabilistic dependencies, i.e. the B_i will correspond to parents of the node H ; therefore, a conditional probability distribution cpd needs to be associated to each clause.

Before continuing to discuss some problems and extensions, let us first illustrate these ideas on our **genetic** example. In doing, so let us slightly redefine the predicates **pc**/2, **mc**/2

²More specifically, it only makes sense to talk about the atoms that are logically entailed by the program, i.e. that are true in all possible Herbrand interpretations, that is the atoms belonging to the so-called least Herbrand model.

```

mc(P)    :- mother(P, M), mc(M), pc(M).
pc(P)    :- father(P, F), mc(F), pc(F).
mc(1mf). mother(1, 1m).
pc(1mf).  father(1, 1f).
mc(1mm).  mother(1m, 1mm).
pc(1mm).  father(1m, 1mf).
mc(1f).
pc(1f).
pt(P)    :- mc(P), pc(P).

```

Figure 8: The truncated **genetic** program.




Figure 9: The structure of the Bayesian network induced by the truncated **genetic** program, see Figure 8. All nodes (with a non-empty set of parents) for the same predicate share the same associated cpds.

and **pt**/2. We will now turn these two-place predicates into predicates of arity 1, the idea being that **pt**(P) will be (in state) true if and only if the phenotype of the person P is green. The other predicates are dealt with in a similar way. Together with the facts for **mother**/2 and **father**/2 (listed in Figure 3), the truncated facts for **mc**/1 and **pc**/1 result in the truncated **genetic** program as shown in Figure 8. This program then induces a Bayesian network with the graphical structure shown in Figure 9. All nodes (with a non-empty set of parents) for the same predicate (i.e., **mc**/1, **pc**/1 and **pt**/1) share the same local probability model. So, the cpd’s for **mc**(1m) and **mc**(1mm) are the same.

There are some problems with this approach. Let us discuss the most notable ones. First, the Bayesian network obtained is required to be acyclic. This is required because the semantics of cyclic Bayesian networks are not well-defined. Therefore, virtually all approaches require acyclicity or enforce other restrictions that guarantee acyclicity.

Second, the domains of the random variables may not be boolean, i.e., they may have more than two possible values, or even be continuous. There are two ways to realize this. First, one could simply allow the domains of the random variables to be arbitrary, cf. [41; 52]. One side-effect is that boolean or logical random variables get mixed up with

other ones. This may sometimes cause some confusion. Second, one could simply add an extra argument to the atoms, cf. [57; 73]. E.g., in the **genetic** domain, this would correspond to the encoding used in Figure 3. The disadvantage of this approach is that one needs to add extra constraints that would state that the phenotype must be either red or white, cf. [73]. This in turn results in a more complicated inference engine.


Finally, it might be the case that there exist multiple ground clauses with the same head. Indeed, consider the program consisting of the following two clauses:

```
p :- q,r.
p :- s,t.
```

The first clause states that $\mathbf{Pa}(\mathbf{p}) = \{\mathbf{q}, \mathbf{r}\}$, the second one that $\mathbf{Pa}(\mathbf{p}) = \{\mathbf{s}, \mathbf{t}\}$. This problem can also arise with a program consisting of a single clause. Indeed, consider $\mathbf{p}(\mathbf{X}) :- \mathbf{q}(\mathbf{X}, \mathbf{Y})$, for which there may be multiple (true) instantiations of $\mathbf{q}(\mathbf{X}, \mathbf{Y})$ for a single \mathbf{X} . Various solutions which might even be combined have been developed for dealing with this situation. First, one could introduce so-called *combination rules*, cf. [73; 45; 52]. Combination rules are functions that would – in the above case – take the corresponding conditional probability distributions $\mathbf{P}(\mathbf{p} \mid \mathbf{q}, \mathbf{r})$ and $\mathbf{P}(\mathbf{p} \mid \mathbf{s}, \mathbf{t})$ as inputs and would produce the desired $\mathbf{P}(\mathbf{p} \mid \mathbf{q}, \mathbf{r}, \mathbf{s}, \mathbf{t})$ as output. Classical examples of combining rules are *noisy-or* and *noisy-and*. For illustration purposes, consider the clauses

```
cold.          fever :- cold.
flu.           fever :- flu.
malaria.       fever :- malaria.
```

together with *noisy-or* as combining rule. The Bayesian network induced is



where *noisy-or* ($\text{cpd}(\text{fever} :- \text{flu})$, $\text{cpd}(\text{fever} :- \text{cold})$, $\text{cpd}(\text{fever} :- \text{malaria})$) is associated with **fever**, cf. [86, page 444]. Then **fever** is false only if all parents are independently false. The probability of **fever** being true then depends on the parents which are true. More precisely,

$$P(\text{fever} = \text{true}) = 1 - \sum_{\substack{X \in \{\text{cold}, \text{flu}, \text{malaria}\} \\ \text{which is true}}} P(X = \text{true})$$

Second, one could employ an *aggregate function* as done in probabilistic relational models (PRMs) [57; 81]. Aggregate functions are well-known from databases; examples include min, max, sum, average, count, etc. To illustrate their use in the present context, reconsider the clause $\mathbf{p}(\mathbf{X}) :- \mathbf{q}(\mathbf{X}, \mathbf{Y})$ and consider that for $\{\mathbf{X} = \mathbf{a}\}$ there are multiple substitutions, let us say $\{\mathbf{Y} = \mathbf{b}, \mathbf{Y} = \mathbf{c}\}$, such that $\mathbf{q}(\mathbf{a}, \mathbf{Y})$ holds. Then $\{\mathbf{q}(\mathbf{a}, \mathbf{b}), \mathbf{q}(\mathbf{a}, \mathbf{c})\}$ are the two random variables that influence $\mathbf{p}(\mathbf{a})$. Furthermore, these random variables will have values, say v_1 and v_2 . An aggregate function τ could then map this set of values $\{v_1, v_2\}$ onto a single value v . Let us use

the notation $\tau(\mathbf{q}(\mathbf{a}, \mathbf{Y}))$ to denote this value v . The random variable $\mathbf{p}(\mathbf{a})$ thus depends on $\tau(\mathbf{q}(\mathbf{a}, \mathbf{Y}))$, which is of course also a random variable. If multiple atoms appeared in the clause, e.g., $\mathbf{p}(\mathbf{X}) :- \mathbf{q}(\mathbf{X}, \mathbf{Y}), \mathbf{r}(\mathbf{X}, \mathbf{Z})$, then each of the atoms would be aggregated independently. The latter clause thus represents the probabilistic dependence of $\mathbf{p}(\mathbf{X})$ on $\tau_1(\mathbf{q}(\mathbf{X}, \mathbf{Y}))$ and $\tau_2(\mathbf{r}(\mathbf{X}, \mathbf{Z}))$, i.e., the cpd associated with the clauses actually specifies for each value in $\tau_1(\mathbf{q}(\mathbf{X}, \mathbf{Y})) \times \tau_2(\mathbf{r}(\mathbf{X}, \mathbf{Z}))$ a distribution for $\mathbf{p}(\mathbf{X})$. The effect of aggregate functions is that they reduce the information in a natural way. This is advantageous because it reduces the size of the conditional probability distributions in a natural way, but – on the other hand – it may also result in a loss of information [46].

Above we introduced the principles underlying logical extensions of Bayesian networks. Let us now briefly survey some of the key representational frameworks that have been developed around these principles.

As mentioned before, the first approach to directly upgrade Bayesian networks to relational representations, seems due to Peter Haddawy [41]. In his initial approach, the random variables correspond to ground atoms. Furthermore, the need for combining rules was alleviated by requiring that for each random variable, there was at most one instantiation of a clause having the random variable as the head.

In later work, Haddawy now together with Ngo [73], used combination rules to deal with the latter problem, and also, the values of the random variables were now written as an extra argument of the atoms, which necessitates the need to add constraints, cf. above. The language, called probabilistic-logic programs (PLPs), was further extended in the knowledge based model construction tradition by distinguishing between logical and probabilistic conditions in clauses. The former then specified the (logical) conditions under which the probabilistic dependencies held. This increased the expressive power and was certainly useful for specifying probabilistic models by hand, but on the other hand, it also significantly extended and therefore complicated the description language. These complications may also explain why – to the authors' knowledge – no (structural) learning results are known for this formalism. The framework has been applied within medical domains [42; 72].

Building on Ngo and Haddawy's work, Kersting and De Raedt [52] introduced the framework of Bayesian logic programs (BLPs). This framework can be viewed as a simplification of that by Ngo and Haddawy in that BLPs employ a minimal set of primitives such that BLPs have both definite clause logic (i.e. pure Prolog) and Bayesian networks as a direct special case. This in turn facilitates the learning, cf. [53; 51; 54]. Bayesian logic programs again view the atoms as the random variables. More precisely, the atoms of the (least) Herbrand model constitute the set of random variables of the induced Bayesian network. BLPs deal with continuous random variables and employ a simpler form of combination rule.

Finally, there is the quite popular formalism of probabilistic relational models³ (PRMs) [57; 81; 31; 33]. This formalism

³Several other extensions of Bayesian network have been developed [60; 58; 61; 80; 81] by Pfeffer et al. before introducing probabilistic relational models. These extensions have been termed frame-based and object-oriented, in which the relational and logical issues are less direct. Therefore, we base our discussion of probabilistic relational models on

does not use definite clause logic (i.e. pure Prolog) as the underlying representational framework, but rather the entity-relationship model. In PRMs, the idea is that the information about one entity type is stored in one relation. E.g., in the **genetic** illustration, this could be realized using the relation `person(Person, MC, PC, PT)`. So each ground atom can store information about multiple attributes and the dependencies are defined at the level of these attributes. Two types of dependencies are allowed: direct dependencies from other attributes (for the same entity, e.g., the blood type of a person could depend on the two chromosomes of that person), or dependencies via so-called *slot chains*. Slot-chains are binary relations (or binary projections of other relations) that relate the attributes of one entity to that of other ones. E.g., in the **genetic** example, the MC attribute of one person would depend via the slot-chain `mother(Person, Mother)` on the attributes MC and PC of the entity corresponding to the **Mother**. Furthermore, in the basic PRM setting the relations among entities (such as **Mother** or **Father** in the **genetic** example) are assumed to be deterministic and given, cf. [33]. This implies that it would be impossible to specify that the probability is 0.70 that Jef is the father of Mary. Nevertheless, some partial solutions for dealing with probabilistic relations have been developed, cf. [34; 31]. The **genetic** example shows that – at the logical level – PRMs essentially define via single clauses probabilistic dependencies between the attributes of various entities. It also shows that probabilistic dependencies between relations are harder to represent and require special treatment. Within the more logical approaches sketched above, these issues are dealt with in a uniform manner because random variables are represented by ground atoms. Indeed, the more logically oriented version of the **genetic** example, introduced above, shows how the `person/4` relation in PRMs would be split-up (or “normalized”) into atoms that directly correspond to the random variables. On the other hand, the advantages of PRMs are that they (1) possess an elegant graphical representation, see Figure 10, (2) may be more efficient to implement and learn, and (3) can test the acyclicity requirement more easily. Other differences with the above mentioned logical notation include the use of aggregate functions to deal with the multiple instantiations problem and the use of a purely relational language, which does not allow for functors (needed for representing infinite structures), but see [79]. PRMs have been successfully applied to various problems such as relational clustering [98], hypertext classification [37], selectivity estimation within databases [38], and modelling identity uncertainty [76], and led to impressive results within computational biology [95; 93; 94].

Recently, Getoor et al. [31; 35; 32] have introduced *stochastic relational models* (SRMs). At the syntactic level SRMs are almost identical to PRMs but SRMs possess a different semantics, which allow PRMs to answer statistical queries such as “What is the probability that a randomly chosen patient is elderly and has a contact who is a coworker?” The key idea underlying SRMs is that probability distributions are defined and used at an abstract level determined by the database schema. At this level, one is not interested in specific entities (such as specific patients, contacts or co-workers) but rather of their properties (such as the probability that a randomly chosen patient is elderly, or a

the presentation by Getoor et al. in [23].




Figure 10: Graphical representation of a probabilistic relational model of the **genetic** domain.

randomly chosen contact is a coworker). Whereas PRMs define dependencies among different entities via slot-chains, SRMs define them via chains of joins. Many of the other issues SRMs are quite similar to those in PRMs, which explains why we will not provide any further details on SRMs in this paper. For more information, see [31; 35; 32], who present efficient algorithms to construct the desired Bayesian network, to estimate (in particular the join) parameters and to learn the dependency structure from a given database. SRMs have also been used for selectivity estimation within databases [31].

4.2 Probabilistic Proofs

To define probabilities on proofs, there are at least two different options. They correspond to the approaches taken in stochastic logic programs [24; 66; 12; 13; 15] and PRISMs respectively [88; 90; 50; 89].

Let us start by introducing stochastic logic programs (SLPs). Stochastic logic programs are a direct upgrade of stochastic context free grammars. This is realized by associating to each clause a probability value. Furthermore, as in the case of stochastic context free grammars, the sum⁴ of the probabilities associated to clauses of the same predicate equals 1.0. To illustrate stochastic logic programs, reconsider the derivation of “*Rice flies like sand*” but now with the definite clause grammar (with the same associated probabilities). Both give the same probabilities to the same (ground) queries. However, stochastic context free grammars constitute very simple stochastic logic programs. Failures are only due to non-matching terminals. In general SLPs, (intensional) clauses may cause already failures. Consider the **genetic** example in Figure 3 and assume that we have associated uniform probabilities to all clauses for the same predicate. Under these assumption one could e.g. show that the probability of the single proof of $\text{mc}(\text{im}, \mathbf{r})$ is $0.2 \cdot 0.5 \cdot 0.2 \cdot 0.2 = 0.004$. However, computing the probabilities of all other ground atoms over `mc/2` shows that they do not sum to 1.0 but to 0.624 (see below) because some of the refutations fail already at the clause level. Stochastic logic programs account for failures by normalization. The normalization constant of an atom over some predicate `mc/2` is the sum of probabilities of all refutations of the most gen-

⁴More general definitions of stochastic logic programs where e.g. the probability values do not sum to 1.0 are discussed by Cussens [12; 14]. Because learning of stochastic logic programs is considered only for ‘normalized’ programs we restrict our attention to them.

eral goal over $mc(X, Y)$, i.e. the sum of the probabilities of $mc(1mf, r)$, $mc(1mm, r)$, $mc(1f, r)$, $mc(1r, r)$, and of $mc(1m, r)$:

$$\begin{aligned}
& 0.2 \\
& + 0.2 \\
& + 0.2 \\
& + 2 \cdot (0.2 \cdot 0.5 \cdot 0.2 \cdot 0.2) \\
& + 2 \cdot [0.5 \cdot 4 \cdot (0.2 \cdot 0.5 \cdot 0.2 \cdot 0.2)] \\
& = 0.624
\end{aligned}$$

Thus, the normalized probability of $\text{:- } mc(1m, r)$ is $(2 \cdot 0.004)/0.624 = 0.012820513$. More details on inference (including approximative inference) can be found in [13; 15]. Stochastic logic programs have been used to specify declarative priors for probabilistic models [3].

Whereas stochastic logic programs attach probability labels to all clauses, PRISMs and the earlier representation by David Poole [82] attach probability labels to facts. Indeed, let a logic program consist of the set of facts E (this could be considered the extensional part) and the set of proper clauses I (the intension). The frameworks by Poole and Sato would then associate a probability label to each of the facts in E . E.g., in the **genetic** example (where we take the truncated encoding of Figure 8), one could associate a probability of 1.0 to the facts for **mother** and **father** and in addition there could be probabilities associated to the facts for **mc/1** and **pc/1**, e.g.,

- (1.0) $mc(1mf)$.
- (1.0) $pc(1mf)$.
- (0.7) $mc(1mm)$.
- (0.7) $pc(1mm)$.
- (1.0) $mc(1f)$.
- (1.0) $pc(1f)$.

A labelled fact $p : f$ then has the meaning that there is a probability of p that the fact f is true. This induces a probability at the level of proofs. Indeed, consider now the atom $pt(1mm)$. The proof of $pt(1mm)$ only succeeds when both $mc(1mm)$ and $pc(1mm)$ are true. These two facts are only true with a probability of $0.7 \cdot 0.7 = 0.49$. Therefore, the probability that the proof of $pt(1mm)$ succeeds is 0.49. Because probability values are associated with facts, no normalization is needed as opposed to stochastic logic programs.

At this point the reader has probably observed that the probability distribution for proofs induces a probability distribution for the other atoms and even for the interpretations. Indeed, because there is only one possible proof for $pt(1mm)$, the probability for this atom to be true is also 0.49. If there would be more proofs, then one would have to sum the probabilities. This in turn defines a probability distribution at the level of interpretations and so the formalisms by Sato and Poole, can also be interpreted from a model-theoretic perspective. Sato has called this a *distributional semantics*, cf. [88], stressing the declarative character of PRISMs. However, the key difference with the Bayesian network approaches presented above is that the former approaches rely more on *logical* inference methods, whereas the latter rely more on *probabilistic* inference mechanisms. In a sense, the assumptions (on which the proofs rely) are probabilistic, whereas in Bayesian networks the inference mechanism is. This in turn explains the strong link with




Figure 11: The graph structure of a Markov model for web navigation.

abductive inference methods, cf. [82]. Finally, let us remark that we discussed simplified versions of Poole's and Sato's frameworks. Originally, they have been introduced to deal with situations where some of the facts are mutually exclusive. E.g., if three genes would be involved (say a , b and c), then the corresponding facts for $mc/2$ and $pc/2$ would be mutually exclusive. This could be specified using logical constraints such as

```

false :- mc(X, a), mc(X, b).
false :- mc(X, a), mc(X, c).
false :- mc(X, b), mc(X, c).
mc(X, a); mc(X, b); mc(X, c) :- person(X).

```

A more compact notation used by Sato and Poole is $\text{disjoint}[p_1: mc(X, a), p_2: mc(X, b), p_3: mc(X, c)]$ where p_i denotes a probability value.

4.3 Intermediate Representations

The above introduced frameworks integrate probabilistic models with very expressive relational representations or even logic programs. This however comes at a computational cost. Therefore, some recent approaches such as relational Markov models (RMMs) [2], hidden tree Markov models (HTMMs) [28; 21], and logical hidden Markov models (LOHMMs) [55] have tried to upgrade some well understood probabilistic representations, such as (Hidden) Markov Models [84]. These approaches can also be viewed as downgrading one of the more expressive probabilistic logics presented above. Whereas the resulting approaches are typically less expressive than the above introduced ones, the advantages are that (1) they are closely related to the underlying representations, (2) they are – in principle – more efficient, and (3) that the learning algorithms for the underlying representations can almost directly be applied.

Let us illustrate this idea on Markov models (where we follow the ideas underlying RMMs and LOHMMs using an example by Anderson et al. [2]). A Markov model is essentially a finite state automaton with probabilities associated to each of the edges instead of symbols from an alphabet. Alternatively, it can be regarded as stochastic regular grammar⁵. A (toy) example for web navigation (within a given academic site) is given in Figure 11. The model denotes the probabilities that if one is at a web page of type X the next web page will be of type Y . E.g., in the graph listed above, the probability of moving from **department** to **lecturer** is 0.3. It should be clear that this model is too simple to be useful for web user modelling because the abstraction level is too high. One alternative would be to build a model with

⁵This is a special case of a stochastic context free grammar in which all rules have exactly one symbol on the right hand.

one node for each of the web pages. This model would however be so huge that it would hardly be usable (let alone learnable). Therefore, a better way, supported by relational or logical Markov models is to use proper relational or logical *atoms* to represent the state instead of propositional symbols. A sequence of possible states could then be e.g.

$$\text{dept}(\text{cs}) \rightarrow \text{course}(\text{cs}, \text{dm}) \rightarrow \text{lecturer}(\text{cs}, \text{pedro}) \rightarrow \text{course}(\text{cs}, \text{stats}) \rightarrow \dots$$

The sequence of pages represented would include the web page of the department of **cs**, the **dm** course, its lecturer **pedro** and the course **stats** taught by the same lecturer. This type of relational and logical sequence could be modelled using the following type of (grammar) rules

(0.7)	dept(D)	→	course(D, C).
(0.2)	dept(D)	→	lecturer(D, L).
			...
(0.3)	course(D, C)	→	lecturer(D, L).
(0.3)	course(D, C)	→	dept(D).
(0.3)	courses(D, C)	→	course(D', C').
			...
(0.1)	lecturer(X, L)	→	course(X, C).
			...

Now, starting from any ground state, e.g., `dept(cs)`, one can determine the possible transitions as well as their probabilities. However, the above logical Markov model is only specifying the probability of a transition to an *abstract state* (i.e., a state that is not ground - it contains variables); in our example, the abstract states are `course(cs, C)` and `lecturer(cs, L)`. To be able to determine the corresponding real states, one also needs to know (1) the domains of the various arguments (in our examples, the set of all lecturers and the set of all courses, say $\text{dom}(L)$ and $\text{dom}(C)$), and (2) a probability distribution on these domains (i.e., \mathbf{P}_L and \mathbf{P}_C) that specifies the probability of selecting a particular instance from these domains, e.g. the probability of selecting `pedro` as a lecturer would be given by $P_L(\text{pedro})$. Given these domains and their corresponding distributions, one can now instantiate the abstract transitions. E.g., starting from `dept(cs)` there would be a probability of $0.7 \cdot P_L(\text{pedro})$ of going to `lecturer(cs, pedro)`. This illustrates the key ideas underlying these representations: proof steps correspond to time steps. Nevertheless, various differences among these approaches exist:

- Relational Markov Models [2] do not allow for variables and unification. Instead they employ a taxonomy to define the domains and probability estimation trees [83] to specify the domain distributions. Anderson et al. [83] present impressive experimental results on adaptive web navigation.
- Logical Hidden Markov Models [55] do allow for variables and unification and they also upgrade hidden Markov models (instead of the simpler Markov models presented above). Kersting et al. [55] present a bioinformatics application on predicting the fold class of proteins.
- Hidden tree Markov models [21] and other hidden Markov models over data structures [28] focus on modelling probability distributions defined over spaces of

trees and graphs. They do not use logical concepts such as predicates. Instead, they adapt automata-theoretical concepts such as deterministic transduction to generically define the dependency structure among random variables. Diligenti et al. present an application on document image classification in [21].

Further possible choices for these approaches exist. One of these is concerned with the level at which to specify the domain probabilities. One can specify them locally, i.e., for each "predicate" independently, or globally. Secondly, how should one combine these probability distributions to account for larger substitutions, e.g., $X = \text{cs}, Y = \text{pedro}$ in $\text{lecturer}(X, Y)$? Either, one could take a naïve approach by assuming that the domains are independent (similar to naïve Bayes), i.e., $P(X = \text{cs}, Y = \text{pedro}) = P(X = \text{cs}) \cdot P(Y = \text{pedro})$, or employ more advanced models such as Bayesian networks or other probabilistic models. More experience and research is necessary in order to fully understand and appreciate the different options, opportunities and limitations of these recent intermediate approaches.

Finally, let us remark that the intermediate representations presented in this section should be viewed as belonging to the *proof-theoretic* approach as probability distributions are defined at the level of derivations and (partial) proofs respectively.

5. LEARNING PROBABILISTIC LOGICS

Learning denotes the process by which one adapts (probabilistic) models on the basis of data. When learning within probabilistic logics (as well as within the more traditional probabilistic representations), one typically distinguishes two tasks. The first task is concerned with the probabilistic part, i.e., the question *where do the numbers come from*; it is the so-called *parameter estimation* problem. In this task, it is assumed that the structure (i.e., the underlying logic program) of the representation is fixed and known, and the parameters (i.e., the probability labels) have to be estimated. In the second task, the so-called *structure learning* or *model selection*, one assumes that neither the logic program nor the parameters are fixed but have to be learned.

Again, it will turn out to be useful to distinguish the model-based from the proof-based representations. It will turn out that — from the inductive logic programming perspective — the distinctions are akin to those between *learning from entailment* and *learning from interpretations*, cf. [17]. We will therefore start this section by discussing the types of data employed for learning. Because one also needs to estimate the parameters when learning the structure of a probabilistic model, we will introduce parameter estimation before structure learning.

5.1 The Evidence

The model-based and the proof-based representations employ different forms of data which we will now sketch.

5.1.1 Model Theoretic

An example or data case used for learning in a Bayesian network setting consists of the random variables specified in the network and a possibly partial assignment of values to states⁶. To illustrate this, reconsider our earlier **alarm**

⁶In this paper, we will assume that all random variables are

example. A possible example for this network could be

$\{\text{burglary} = \text{false}, \text{earthquake} = \text{true}, \text{alarm} = ?$
 $\text{johncalls} = ?, \text{marycalls} = \text{true}\}$

This example partially describes a particular situation that arose in **alarm**. Judea was called by Mary after an earthquake took place. Judea neither knew whether the alarm went off nor whether John tried to call him. Therefore, when working with logical or relational upgrades of Bayesian networks, it is natural to work with examples in the form of interpretations. Indeed, the atoms in the interpretations specify the random variables, and their values specify the state they are in. E.g., in the **genetic** illustration, one could have the following example:

$\{\text{pc}(1\text{f}) = \text{true}, \text{mc}(1\text{f}) = \text{true}, \text{pc}(1\text{m}) = \text{true}, \text{mc}(1\text{m}) = ?$,
 $\text{father}(1, 1\text{f}) = \text{true}, \text{mother}(1, 1\text{m}) = \text{true}, \text{pc}(1) = ?$,
 $\text{mc}(1) = ?, \text{pt}(1) = \text{true}\}$.

Remark that — from a logical perspective — there are two constraints on this type of examples. First, the random variables specified by a single example can be considered a Herbrand interpretation, and this interpretation should be a model of the unknown target program. This is akin to the *learning from interpretations* setting from the field of inductive logic programming. Second, the Bayesian network structure induced by the target program for this particular example should be acyclic. These two properties will turn out to be important when learning the structure.

5.1.2 Proof Theoretic

To obtain insight into the nature of the examples for the proof theoretic setting, it is a good idea to look at the nature of the data when learning stochastic context free grammars. When learning stochastic context free grammars, the examples given are strings that should be accepted (with an unknown probability) by the target grammar. E.g., when learning a stochastic context free grammar in a natural language domain for English, the examples would be English sentences. Furthermore, these sentences should be derivable by the target grammar. When working with definite clause grammars (as in Figure 4 (a)) sentences would be represented by facts. E.g., the sentence for the non-terminal symbol **S** would be represented as $\text{s}([\text{the}, \text{dog}, \text{barks}], [])$. Thus, within the proof oriented setting of probabilistic logics, the evidence will correspond to facts (or even clauses). Furthermore, these examples e should be logically entailed by the target program T , i.e., $T \models e$. This is akin to the traditional *learning from entailment* setting in inductive logic programming, cf. [69].

5.1.3 Intermediate Representations

Intermediate representations, such as (hidden) Markov models, are typically learned from examples that correspond to sequences of states that arose in particular situations. E.g., the Markov model for the web site domain, could be learned from examples such as **department**, **lecturer**, **course**, **lecturer**, ..., or in the first order case by sequences such as the one specified in Sec-

tion 4.3: known, i.e., that no *latent* variables, whose states are never observed have to be introduced by the learning algorithm.

tion 4.3:

$\{\text{dept}(\text{cs}) \rightarrow \text{course}(\text{cs}, \text{dm}) \rightarrow \text{lecturer}(\text{cs}, \text{pedro}) \rightarrow$
 $\text{course}(\text{cs}, \text{stats}) \rightarrow \dots\}$

Such sequences correspond to a kind of (partial) *trace* (or derivation) of the underlying proof system. Learning from traces has also been considered in the field of inductive logic programming, see [96; 5] for two well-known examples where traces have been used to induce logic programs. Traces impose a strong logical constraint on the underlying program. Indeed, for each two subsequent states s_1 and s_2 in a trace (for a logical Markov model), it is required that there exists a clause $b \rightarrow h$ and a substitution θ such that $s_1 = b\theta$ and $s_2 = h\theta$.

5.2 Parameter Estimation

Parameter estimation methods start from a set of examples E and a given logic program L of a probabilistic logical model, and aims at computing the values of the parameters λ^* that best explain the data. So, λ is a set of parameters (i.e. the quantitative part of the model) and can be represented as a vector. To measure the extent to which a model fits the data, one usually employs the likelihood of the data. The likelihood of the data is the probability of the observed data as a function of the unknown parameters with respect to the current model, i.e. $P(E|L, \lambda)$. Thus, maximum likelihood estimation (MLE) aims at finding $\lambda^* = \text{argmax}_{\lambda} P(E|L, \lambda)$. As $\text{argmax}_{\lambda} P(E|L, \lambda) = \text{argmax}_{\lambda} \log P(E|L, \lambda)$, one often works with the log-likelihood $\log P(E|L, \lambda)$, which is easier to manipulate. Often, one assumes that the examples are independently sampled from identical distributions, i.e. $\log P(E|L, \lambda) = \sum_{e \in E} \log P(e|L, \lambda)$. Variants or extensions of the MLE criterion presented are sometimes used as well. They include: a Bayesian approach which would take into account parameter priors or the minimum description length (MDL) principle that is often used for structure learning, cf. Section 5.3.

To illustrate MLE, let us consider tossing a coin. The coin can either be heads or tails. Assume that the evidence is

$E = \{\{\text{toss} = \text{head}\}, \{\text{toss} = \text{tail}\}, \{\text{toss} = \text{tail}\}\}$

(respectively $\{\{\text{head}\}, \{\text{tail}\}, \{\text{tail}\}\}$ in the proof-based notation), and that our current quantitative part is $P(\text{toss} = \text{head}) = p$. Then

$$\begin{aligned} \log P(E|L, \lambda) &= \sum_{e \in E} \log P(e|L, \lambda) \\ &= 1 \cdot \log p + 2 \cdot \log(1 - p) \end{aligned}$$

We can maximize the log-likelihood by setting the partial derivative w.r.t. to p equal to 0.0:

$$\frac{\partial \log P(E|L, \lambda)}{\partial p} = \frac{1}{p} - \frac{2}{1 - p} = 0$$

Solving this give $p = \frac{1}{3}$, which is the “intuitive” estimate for p namely the proportion of heads which have been seen in the examples. More formally,

$$p = \frac{\text{count}(\text{toss} = \text{head})}{\text{count}(\text{toss} = \text{head}) + \text{count}(\text{toss} = \text{tail})}$$

where $\text{count}(\cdot)$ denotes the frequency of an event. Thus, MLE reduces to frequency counting.

However, in the presence of missing data, the maximum likelihood estimate typically cannot be written in closed form. It is a numerical optimization problem, and all known algorithms involve nonlinear optimization. The most commonly adapted technique for probabilistic logic learning is the Expectation-Maximization (EM) algorithm [20; 65]⁷. EM is based on the observation that learning would be easy (i.e., correspond to frequency counting), if the values of all the random variables were known. Therefore, it first estimates these values, then uses these to maximize the likelihood, and then iterates. More specifically, EM assumes that the parameters have been initialized (e.g., at random) and then iteratively perform the following two steps until convergence⁸:

- **E-Step:** On the basis of the observed data and the present parameters of the model, compute a distribution over all possible completions of each partially observed data case.
- **M-Step:** Using each completion as a fully-observed data case weighted by its probability, compute the updated parameter values using (weighted) frequency counting.

The frequencies over the completions are called the *expected counts*. Let us illustrate EM using the evidence

$$E = \{\{\mathbf{toss} = \mathit{head}\}, \{\mathbf{toss} = ?\}, \{\mathbf{toss} = \mathit{tail}\}\}$$

and assume that in the current model $P(\mathbf{toss} = \mathit{head}) = p$. The completions of the data case then are

$$\begin{aligned} 1 : \{\mathbf{toss} = \mathit{head}\}, 1 : \{\mathbf{toss} = \mathit{tail}\}, \\ p : \{\mathbf{toss} = \mathit{head}\}, 1 - p : \{\mathbf{toss} = \mathit{tail}\}. \end{aligned}$$

The updated maximum likelihood estimation of the probability that the coin will be head is then $\frac{1+p}{3}$. Assuming $p = 0.8$, this gives 0.6, and iterating yields 0.53, 0.511, ... which converges to 0.5.

Let us now discuss how EM has been applied within the various probabilistic logics introduced earlier. For the sake of simplicity, we will only state the key ideas and will not address any advanced issues such as efficiency concerns.

5.2.1 Model Theoretic

Parameter estimations for probabilistic-logic programs [59], probabilistic relational models [30; 31; 33], and Bayesian logic programs [51; 54] all follow the same principle: The given data and the current model induce a Bayesian network explaining each data cases. Then, the parameters of the induced Bayesian network are estimated using standard Bayesian network parameter estimation methods, such as those discussed in Heckerman's tutorial [44]. For instance the examples in Section 5.1.1 together with the **genetic** program would induce the Bayesian network shown in Figure 9. All nodes of the Bayesian network, which do not occur in the evidence but are introduced by the program,

⁷Some approaches employ gradient-based approaches instead of EM, e.g., Muggleton [67; 68] for structural learning of SLPs, and Kersting and De Raedt [51; 54] developed both EM-based and gradient-based approaches for parameter estimation of BLPs.

⁸This is only the underlying idea of the EM. More formally, the E-step consists of computing the expectation of the likelihood given the old parameters and the observed data. Then, the M-step consists of maximizing the expected likelihood w.r.t. the parameters.

are not observed, i.e. they are assigned a question mark. Thus the evidence corresponds to a partial joint state of the resulting Bayesian network. It is completed by the E-step, which relies on traditional Bayesian network inference and which is used to determine the distribution of the values for the unobserved states. The improved estimates of the parameters in a node (i.e., the values in the cpds) are then computed in the M-step, by dividing the expected counts of corresponding node-parents and parents joint states. The only difference with standard Bayesian network parameter estimation is that parameters for different nodes in the network — those corresponding to different ground instances of the same clause — are forced to be identical. This technique is akin to that in recurrent neural network [99] or dynamic Bayesian networks [19]. It however requires a unique assignment of parent nodes to clauses. Aggregate functions guarantee this for PRMs as the aggregated examples constitute the states of the nodes in the induced Bayesian network. Within PLPs and BLPs, uniqueness is enforced by using decomposable combining rules [59; 51]. The effect of a decomposable combining rule can be represented using extra nodes in the induced Bayesian network. Most combining rules commonly employed in Bayesian networks such as *noisy or*, *noisy and*, or linear regression are decomposable.

5.2.2 Proof Theoretic

Whereas model-based approaches complete the data in terms of a Bayesian network, proof-theoretic approaches complete the data based on refutations and failures.

For stochastic logic programs, Cussens [14] introduced the *failure-adjusted maximization* (FAM) algorithm. In FAM, the logical part of the SLP is fixed and given, and the parameters have to be learned on the basis of the evidence. The evidence or the examples are atoms for a predicate p/m . They are assumed to be generated from the target SLP according to the probability distribution that it induces over p/m . This implies that the examples are already logically entailed by the SLP. The parameters are estimated by computing for each example the SLD tree and treating each path from root to leaf in the SLD tree as one possible completion. For instance the evidence $\mathbf{pt}(\mathbf{1mm}, \mathbf{red})$ yields five possible completions, namely the five paths of the lower SLD tree in Figure 5, of which four are failed and one constitutes a proof⁹. Secondly, the completions are weighted with the product of probabilities presently associated with clauses used in the corresponding derivations, cf. Section 4.2. Finally, the improved estimates for each clause are obtained by dividing the clause's expected counts by the sum of the expected counts of clauses for the same predicate.

The EM algorithm for PRISM programs [50; 91; 49; 92] is similar in spirit. The main differences are that (1) no probability values are associated to (intensional) clauses, (2) failed derivations have a probability of zero, and (3) the **disjoint** representation introduced in Section 4.2 is assumed. Given an example, all explanations S of the example are computed. Similar to abductive reasoning, an explanation of an example is a set of facts S that allows the example to be proven. For instance, $\{\mathbf{pc}(\mathbf{1mm}, \mathbf{r})\}$ is the only explanation for $\mathbf{pt}(\mathbf{1mm}, \mathbf{red})$, i.e. the failure derivation in Figure 5 is neglected. Thus, the explana-

⁹The approach is called *failure-adjusted maximization* because both failed and successful derivations are taken into account.

tions constitute the completions of the data, and they are weighted with their probability $P(S)$ as specified by the current parameters. The maximum likelihood estimation of the probability value associated with a fact, say $\text{pc}(1\text{mm}, r)$, is then the fact's expected count over all explanations where it is true divided by the expected counts of all corresponding mutually exclusive facts, e.g. given the declaration $\text{disjoint}[p_1: \text{pc}(1\text{mm}, r), p_2: \text{pc}(1\text{mm}, g)]$, $\text{pc}(1\text{mm}, r)$ and $\text{pc}(1\text{mm}, g)$ are mutually exclusive. To speed up computations, Sato and Kameya employ well-known tabulation techniques from the field of computational logic.

As already argued above, the intermediate representations, such as LOHMMs and RMMs can be viewed as belonging to the *proof-theoretic* approach. Therefore, we will omit a detailed discussion of parameter estimation for these methods.

5.3 Structure Learning

In structure learning, one starts from a set of examples E , a language bias \mathcal{L} that determines the set of possible hypotheses, and aims at computing a hypothesis $H^* \in \mathcal{L}$ such that

1. H^* logically covers the examples E , i.e., $\text{cover}(H^*, E)$, and
2. the hypothesis H^* is optimal w.r.t. some scoring function score, i.e., $H^* = \text{argmax}_{H \in \mathcal{L}} \text{score}(H, E)$.

The hypotheses H are of the form (L, λ) where L is the logical part L and λ the vector of parameter values as in Section 5.2. The coverage criterion used will typically depend on the type of data and representation formalisms considered, cf. Section 5.1, where we have discussed three different types of data for our three types of representation formalisms. In addition, as mentioned in Section 5.2, various scoring functions can be employed as well. A simple example would be $\text{score}(H, E) = P(E \mid H) = P(E \mid L, \lambda)$, i.e., the likelihood criterion. The type of data, the coverage criterion, the language bias and the scoring function are the central components of any structure learning method. However, in addition, various enhancements or variants can be considered as well. These include a possible initial hypothesis that could serve as a starting point for guiding the search and the incorporation of background knowledge. Background knowledge could consist of a set of given and fixed clauses (possibly with attached probability labels) that would provide useful information about the problem domain. The use of an initial hypothesis is common in many probabilistic learning as well as in theory revision approaches [100] to inductive logic programming. Background knowledge is typically employed in multi-relational data mining and inductive logic programming, cf. [69]. Further distinctions could be made according to whether the data are given initially and fixed or whether they are gathered incrementally. Such distinctions and enhancements are quite similar to those made in logic learning and probabilistic learning. Therefore, we will not provide more details on these aspects as doing so would lead us too far. Instead, we will concentrate on the generalized problem sketched above.

Nearly all (score-based) approaches to structure learning perform a heuristic search through the space of possible hypotheses as determined by \mathcal{L} . Typically, hill-climbing or beam-search is applied until the candidate hypothesis




Figure 12: Refinement operators during structure learning of Bayesian networks. We can add a proposition to the body of a clause (respectively fact) or delete it from the body.




Figure 13: First-order refinement operators during structure learning. Atoms can be added to or deleted from the body of a clause. Facts can be added to and deleted from the program.

(1) satisfies $\text{cover}(H, E)$ and (2) the $\text{score}(H, E)$ is no longer improving. The steps in the search-space are made through the application of so-called refinement operators [96; 74], which make perform small modifications to a hypothesis. From a logical perspective, these refinement operators typically realize elementary generalization and specialization steps (usually, under θ -subsumption). In probabilistic learning, certainly when working with graphical models, larger steps are sometimes taken, see also below. Let us now illustrate how these refinement operators work. Assume the current candidate hypothesis H consists of the set of clauses $\{c_1, \dots, c_n\}$. Then the following results of applying a (specialization) refinement operator ρ on the overall hypothesis H can be imagined. Notice that ρ operates at the hypothesis or theory level and also that it employs a refinement operator ρ_c at clause level¹⁰:

- $\rho(H) = (H \setminus \{c_i\}) \cup \rho_c(c_i)$ where $c_i \in H$.
- $\rho(H) = H \cup \{f\}$ where f is a fact with $f \notin H$.
- $\rho_c(h \text{ :- } b_1, \dots, b_m) = \{h \text{ :- } b_1, \dots, b_m, b \mid b \text{ is an atom}\}.$

In addition, one typically also employs a dual generalization operator. This operator would generalize hypotheses by deleting atoms in clauses and facts from hypotheses. The refinement operator is illustrated in Figure 12 for propositional clauses and in Figure 13 for first order ones.

By now, we are able to describe some of the key approaches to structure learning in probabilistic logics.

Bayesian networks directly correspond to propositional clauses. In the typical structure learning approach, both a propositional generalization and specialization operator are applied. In one step, atoms are either added to or deleted from clauses. Furthermore, there is exactly one

¹⁰At this point, note that we employ a simplified and idealized refinement operator. Some essential features, such as substitutions are not taken into account for reasons of simplicity, cf. [69; 74].

clause for each of proposition. Special care is taken that the resulting network is acyclic [44]. The state-of-the-art learning algorithm is the structural EM (SEM) [29]. It adapts the standard EM algorithm for structure learning. The key idea is that the expected counts are not computed anew for every structure that is proposed, but only after several iterations. This leads to an improved efficiency.

For PRMs, SRMs and BLPs, each application of a refinement operator at the hypothesis level acts as a kind of macro at the level of the induced Bayesian network. Indeed, when a literal is added to the logical hypothesis, this corresponds to adding multiple edges to the induced Bayesian network on the data. Consequently, Bayesian network learning techniques such as the SEM have been adopted. **Probabilistic relational models** can be represented by clauses such as $p(\mathbf{X}) : - \tau_1(q(\mathbf{X}, \mathbf{Y})), \tau_2(r(\mathbf{X}, \mathbf{Z}))$, where the τ_i denote different aggregate functions and where there is at most one clause defining a predicate (such as $p/1$). Structure learning with probabilistic relational models now occurs by refining this type of clauses [30; 36; 34; 33; 31]¹¹. Aggregated literals such as $\tau_1(q(\mathbf{X}, \mathbf{Y}))$ can be added to or deleted from clauses. At this point, the graphical notation for probabilistic relational models (illustrated in Figure 10) is convenient for visualizing this process. Adding or deleting aggregated literals to or from a clause directly corresponds to the same operation on the arcs. This in turn clarifies the connection to traditional Bayesian network learning. As in Bayesian networks, special care is taken to guarantee that regardless of the extension used the probabilistic relational model will be acyclic. **Bayesian logic programs** are represented by regular clauses, on which the typical refinement operators from inductive logic programming can be applied [53; 54]. However, in Bayesian logic programs, one must take into account the covers constraint. More formally, it is required that the examples are models of the Bayesian logic programs, i.e. $cover(H, e) = true$ if and only if e is model of H . This requirement is needed because — as argued in Section 4.1 — the set of random variables defined by a Bayesian logic program corresponds to a Herbrand model. This is akin to the *learning from interpretations* setting in inductive logic programming [18]. The requirement is enforced when learning the structure of Bayesian logic programs by starting from an initial Bayesian logic programs that satisfies this requirement (such a hypothesis can be computed using the clausal discovery engine in [18]) and from then on only considering refinements that do not result in a violation. In addition, acyclicity is enforced by checking for each refinement and each example that the induced Bayesian network is acyclic.

The framework of stochastic logic programs is a representative of the proof-theoretic probabilistic logics. It differs from the frameworks previously discussed in this section in that the learning problem is not related to that of Bayesian networks. Nevertheless, **stochastic logic programs** are similar to Bayesian logic programs in that they are represented by regular clauses on which the typical refinement operators apply. However, whereas Bayesian logic programs are learned from interpretations,

stochastic logic programs are learned from entailment. Indeed, the examples in stochastic logic programs are facts and these facts must be logically entailed by the target (stochastic) logic program. Thus the coverage requirement states that $H \models e$ for stochastic logic program. Solving the general structure learning problem for stochastic logic programs thus involves applying a refinement operator at the *theory* level (i.e. considering multiple predicates) under entailment. This problem has been studied within the field of inductive logic programming under the term theory revision [96; 16; 100] and is known as a very hard problem. This may explain why this most general form of structure learning in stochastic logic programs (and the related PRISMs) has not yet been addressed. So far, the only contributions to structure learning of stochastic logic programs are restricted to learning missing clauses for a single predicate. In [70; 67], Muggleton introduced a two-phase approach that separates the structure learning aspects from the parameter estimation phase. In a more recent approach [68] Muggleton presents an initial attempt to integrate both phases for single predicate learning.

To learn models within the intermediate representations, one can take advantage of their simplified underlying logic program. **Logical and relational Markov models** can be represented by transition rules of the form $B \rightarrow H$ where B and H are atoms representing the abstract states. One could — in principle — consider each possible abstract transition $B \rightarrow H$ and estimate the corresponding transition probability directly from the data. However, the number of such transitions is way too large for this approach to work in practice. Therefore, other approaches — that only select the most informative transitions — are needed. For RMMs, this is realized by 1) employing one abstract transition for each predicate pair, and 2) by assigning a *probability estimation tree* [83], a kind of decision tree, to each such predicate. The probability estimation tree will assign to any given concrete state in the body part of the transition a probability distribution over the possible resulting states. So, probability estimation trees encode a kind of abstraction hierarchy. Kersting et al. [56] propose a different approach, which is essentially a structural, generalized EM algorithm using refinement operators for the LOHMM formalism.

6. CONCLUSIONS

An overview and survey of the new and exciting area of probabilistic logic learning has been presented. It combines principles of probabilistic reasoning, logical representations and statistical learning into a coherent whole. The techniques of probabilistic logic learning were analyzed starting from a logical (or inductive logic programming) perspective. This turned out to be quite useful for obtaining an appreciation of the differences and similarities among the various frameworks and formalisms that have been contributed to date. In particular, the distinction between a model- and a proof-theoretic view was used for clarifying the relation among the logical upgrades of Bayesian networks (such as PLPs, PRMs, BLPs, etc.) and grammars (such as PRISMs and SLPs). This distinction is not only relevant at the representational level but also for learning. It turned out that one learns from interpretations in the model-theoretic approaches, from entailment in the proof-theoretic ones, and from traces in the intermediate ones (such as RMMs and LOHMMs). Fur-

¹¹Please note that the original work on PRMs does not present the learning problem in terms of clauses.

thermore, principles of both statistical learning and inductive logic programming (or multi-relational data mining) are employed for learning the parameters and structure of the probabilistic logics considered. The authors hope that this survey provides a useful perspective on probabilistic logic learning that may inspire the reader to contribute to this challenging and exciting research area.

Acknowledgements

This work benefited from the European Union IST project IST-2001-33053 (Application of Probabilistic Inductive Logic Programming – APRIL). The authors thank Lise Getoor for providing the graphical representation of the PRM in Figure 10.

7. REFERENCES

- [1] J. Allen. *Natural Language Understanding*. Benjamin/Cummings Series in Computer Science. Benjamin/Cummings Publishing Company, 1987.
- [2] C. R. Anderson, P. Domingos, and D. S. Weld. Relational Markov Models and their Application to Adaptive Web Navigation. In D. Hand, D. Keim, O. R. Zaïne, and R. Goebel, editors, *Proceedings of the Eighth International Conference on Knowledge Discovery and Data Mining (KDD-02)*, pages 143–152, Edmonton, Canada, 2002. ACM Press.
- [3] N. Angelopoulos and J. Cussens. Markov chain Monte Carlo using tree-based priors on model structure. In J. Breese and D. Koller, editors, *Proceedings of the Seventeenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-01)*, pages 16–23, Seattle, Washington, USA, 2001. Morgan Kaufmann.
- [4] F. Bacchus. Using first-order probability logic for the construction of bayesian networks. In D. Heckerman and A. Mamdani, editors, *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93)*, pages 219–226, Providence, Washington, DC, USA, 1993. Morgan Kaufmann.
- [5] F. Bergadano and D. Gunetti. *Inductive Logic Programming: From Machine Learning to Software Engineering*. MIT Press, 1996.
- [6] J. S. Breese. Construction of Belief and decision networks. *Computational Intelligence*, 8(4):624–647, 1992.
- [7] J. S. Breese, R. P. Goldman, and M. P. Wellman. Introduction to the special section on knowledge-based construction of probabilistic and decision models. *Cybernetics*, 24(11):1577–1579, 1994.
- [8] W. Buntine. A guide to the literature on learning probabilistic networks from data. *IEEE Transaction on Knowledge and Data Engineering*, 8:195 – 210, 1996.
- [9] K. L. Clark and F. G. McCabe. PROLOG: A Language for Implementing Expert Systems. In J. E. Hayes, D. Michie, and Y.-H. Pao, editors, *Machine Intelligence*, volume 10, pages 455–470. Ellis Horwood, Chichester, 1982.
- [10] R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic networks and expert systems*. Springer-Verlag, 1999.
- [11] M. Craven and S. Slattery. Relational Learning with Statistical Predicate Invention: Better Models for Hypertext. *Machine Learning*, 43(1–2):97–119, 2001.
- [12] J. Cussens. Loglinear models for first-order probabilistic reasoning. In K. B. Laskey and H. Prade, editors, *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pages 126–133, Stockholm, Sweden, 1999. Morgan Kaufmann.
- [13] J. Cussens. Stochastic logic programs: Sampling, inference and applications. In C. Boutilier and M. Goldszmidt, editors, *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-00)*, pages 115–122, Stanford, CA,, USA, 2000. Morgan Kaufmann.
- [14] J. Cussens. Parameter estimation in stochastic logic programs. *Machine Learning*, 44(3):245–271, 2001.
- [15] J. Cussens. Statistical aspects of stochastic logic programs. In T. Jaakkola and T. Richardson, editors, *Proceedings of the Eighth International Workshop on Artificial Intelligence and Statistics 2001*, pages 181–186, Key West, Florida, USA, 2001. Morgan Kaufmann.
- [16] L. De Raedt. *Interactive Theory Revision: An Inductive Logic Programming Approach*. Academic Press, 1992.
- [17] L. De Raedt. Logical settings for concept-learning. *Artificial Intelligence*, 95(1):197–201, 1997.
- [18] L. De Raedt and L. Dehaspe. Clausal discovery. *Machine Learning*, 26(2-3):99–146, 1997.
- [19] T. Dean and K. Kanazawa. Probabilistic temporal reasoning. In T. M. Mitchell and R. G. Smith, editors, *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pages 524–528, St. Paul, MN, USA, 1988. AAAI Press / The MIT Press.
- [20] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc.*, B 39:1–39, 1977.
- [21] M. Diligenti, P. Frasconi, and M. Gori. Hidden Tree Markov Models for Document Image Classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(4):519–523, 2003.
- [22] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological Sequence Analysis: Probabilistic models of proteins and nucleic acids*. Cambridge University Press, 1998.
- [23] S. Džeroski and N. Lavrač. *Relational Data Mining*. Springer-Verlag, 2001.
- [24] A. Eisele. Towards probabilistic extensions of constraint-based grammars. In J. Dörne, editor, *Computational Aspects of Constraint-Based Linguistics Description-II*. DYNA-2 deliverable R1.2.B, 1994.

- [25] William Feller. *An Introduction to Probability Theory and its Applications: Volume 1*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, 3rd edition, 1968.
- [26] P. Flach. *Simply logical: intelligent reasoning by example*. John Wiley and Sons Ltd., 1994.
- [27] P. Flach and N. Lachiche. 1BC: A first-order Bayesian classifier. In S. Džeroski and P. Flach, editors, *Proceedings of the Ninth International Workshop on Inductive Logic Programming (ILP-99)*, volume 1634 of *LNAI*, pages 92–103, Bled, Slovenia, 1999. Springer.
- [28] P. Frasconi, M. Gori, and A. Sperduti. A general framework for adaptive processing of data structures. *IEEE Transactions on Neural Networks*, 9(5):768–786, 1998.
- [29] N. Friedman. The Bayesian Structural EM Algorithm. In G. F. Cooper and S. Moral, editors, *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 129–138, Madison, Wisconsin, USA, 1998. Morgan Kaufmann.
- [30] N. Friedman, L. Getoor, D. Koller, and A. Pfeffer. Learning probabilistic relational models. In T. Dean, editor, *Proceedings of the Sixteenth International Joint Conferences on Artificial Intelligence (IJCAI-99)*, pages 1300–1309, Stockholm, Sweden, 1999. Morgan Kaufmann.
- [31] L. Getoor. *Learning Statistical Models from Relational Data*. PhD thesis, Stanford University, 2001.
- [32] L. Getoor, N. Friedman, and D. Koller. Learning Structured Statistical Models from Relational Data. *Linköping Electronic Articles in Computer and Information Science*, 7(13), 2002.
- [33] L. Getoor, N. Friedman, D. Koller, and A. Pfeffer. Learning probabilistic relational models. In S. Džeroski and N. Lavrač, editors, *Relational Data Mining*. Springer-Verlag, 2001.
- [34] L. Getoor, N. Friedman, D. Koller, and B. Taskar. Learning Probabilistic Models of Relational Structure. In C. E. Brodley and A. Pohorecký Danyluk, editors, *Proceedings of the Eighteenth International Conference on Machine Learning (ICML-01)*, pages 170–177, Williamstown, MA, USA, 2001. Morgan Kaufmann.
- [35] L. Getoor, D. Koller, and B. Taskar. Statistical Models for Relational Data. In S. Džeroski and L. De Raedt, editors, *Workshop Notes of the KDD-02 Workshop on Multi-Relational Data Mining (MRDM-02)*, 2002.
- [36] L. Getoor, D. Koller, B. Taskar, and N. Friedman. Learning probabilistic relational models with structural uncertainty. In L. Getoor and D. Jensen, editors, *Proceedings of the AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, pages 13–20, 2000.
- [37] L. Getoor, E. Segal, B. Taskar, and D. Koller. Probabilistic Models of Text and Link Structure for Hypertext Classification. In *Workshop Notes of IJCAI-01 Workshop on ‘Text Learning: Beyond Supervision’*, Washington, USA, 2001.
- [38] L. Getoor, B. Taskar, and D. Koller. Selectivity Estimation using Probabilistic Relational Models. In W. G. Aref, editor, *Proceedings of the ACM Special Interest Group on Management of Data Conference (SIGMOD-01)*, Santa Barbara, CA, USA, 2001.
- [39] S. Glesner and D. Koller. Constructing Flexible Dynamic Belief Networks from First-Order Probabilistic Knowledge Bases. In Ch. Froidevaux and J. Kohlas, editors, *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU-95)*, volume 946 of *LNCS*, pages 217 – 226, Fribourg, Switzerland, 1995. Springer-Verlag.
- [40] R. P. Goldman and E. Charniak. Dynamic construction of belief networks. In P. P. Bonissone, M. Henrion, L. N. Kanal, and J. F. Lemmer, editors, *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence (UAI-90)*, pages 171–184, Cambridge, MA, USA, 1990. Elsevier.
- [41] P. Haddawy. Generating Bayesian networks from probabilistic logic knowledge bases. In R. López de Mántaras and D. Poole, editors, *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-1994)*, pages 262–269, Seattle, Washington, USA, 1994. Morgan Kaufmann.
- [42] P. Haddawy, J. W. Helwig, L. Ngo, and R. A. Krieger. Clinical Simulation using Context-Sensitive Temporal Probability Models. In *Proceedings of the Nineteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-95)*, 1995.
- [43] J. Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46:311–350, 1989.
- [44] D. Heckerman. A Tutorial on Learning with Bayesian Networks. Technical Report MSR-TR-95-06, Microsoft Research, 1995.
- [45] M. Jaeger. Relational Bayesian networks. In D. Geiger and P. P. Shenoy, editors, *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, pages 266–273, Providence, Rhode Island, USA, 1997. Morgan Kaufmann.
- [46] D. Jensen and J. Neville. Linkage and autocorrelation cause feature selection bias in relational learning. In C. Sammut and A. G. Hoffmann, editors, *Proceedings of the Nineteenth International Conference on Machine Learning (ICML-02)*, pages 259–266, Sydney, Australia, 2002. Morgan Kaufmann.
- [47] F. V. Jensen. *Bayesian networks and decision graphs*. Springer-Verlag New, 2001.
- [48] M. I. Jordan, editor. *Learning in Graphical Models*. Kluwer Academic Publishers. Reprinted by MIT Press, 1998.
- [49] Y. Kameya and T. Sato. Efficient EM learning with tabulation for parameterized logic programs. In J. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, *Proceedings of the First International*

- Conference on Computational Logic (CL-00)*, volume 1861 of *LNAI*, pages 269–294. Springer-Verlag, 2000.
- [50] Y. Kameya, N. Ueda, and T. Sato. A graphical method for parameter learning of symbolic-statistical models. In *Proceedings of the Second International Conference on Discovery Science (DS-99)*, volume 1721 of *LNAI*, pages 254–276, Kanagawa, Japan, 1999. Springer-Verlag.
 - [51] K. Kersting and L. De Raedt. Adaptive Bayesian Logic Programs. In C. Rouveirol and M. Sebag, editors, *Proceedings of the Eleventh Conference on Inductive Logic Programming (ILP-01)*, volume 2157 of *LNCS*, Strasbourg, France, 2001. Springer.
 - [52] K. Kersting and L. De Raedt. Bayesian logic programs. Technical Report 151, University of Freiburg, Institute for Computer Science, April 2001. (submitted).
 - [53] K. Kersting and L. De Raedt. Towards Combining Inductive Logic Programming and Bayesian Networks. In C. Rouveirol and M. Sebag, editors, *Proceedings of the Eleventh Conference on Inductive Logic Programming (ILP-01)*, volume 2157 of *LNCS*, Strasbourg, France, 2001. Springer.
 - [54] K. Kersting and L. De Raedt. Principles of Learning Bayesian Logic Programs. Technical Report 174, University of Freiburg, Institute for Computer Science, June 2002. (submitted).
 - [55] K. Kersting, T. Raiko, S. Kramer, and L. De Raedt. Towards discovering structural signatures of protein folds based on logical hidden markov models. In R. B. Altman, A. K. Dunker, L. Hunter, T. A. Jung, and T. E. Klein, editors, *Proceedings of the Pacific Symposium on Biocomputing*, pages 192 – 203, Kauai, Hawaii, USA, 2003. World Scientific.
 - [56] K. Kersting, T. Raiko, and L. De Raedt. A Structural GEM for Learning Logical Hidden Markov Models. In S. Džeroski, L. De Raedt, and S. Wrobel, editors, *Workshop Notes of the KDD-03 Workshop on Multi-Relational Data Mining (MRDM-03)*, 2003. (to appear).
 - [57] D. Koller. Probabilistic relational models. In S. Džeroski and P. Flach, editors, *Proceedings of Ninth International Workshop on Inductive Logic Programming (ILP-99)*, volume 1634 of *LNAI*, pages 3–13, Bled, Slovenia, 1999. Springer.
 - [58] D. Koller, A. Levy, and A. Pfeffer. P-classic: A tractable probabilistic description logic. In *Proceedings of the Fourteenth National Conference on AI*, pages 390–397, Providence, Rhode Island, August 1997.
 - [59] D. Koller and A. Pfeffer. Learning probabilities for noisy first-order rules. In *Proceedings of the Fifteenth Joint Conference on Artificial Intelligence (IJCAI-97)*, pages 1316–1321, Nagoya, Japan, 1997.
 - [60] D. Koller and A. Pfeffer. Object-oriented Bayesian networks. In D. Geiger and P. P. Shenoy, editors, *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, pages 302–313, Providence, Rhode Island, USA, 1997. Morgan Kaufmann.
 - [61] D. Koller and A. Pfeffer. Probabilistic frame-based systems. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 580–587, Madison, Wisconsin, USA, July 1998. AAAI Press.
 - [62] N. Lachiche and P. Flach. 1BC2: A True First-Order Bayesian Classifier. In S. Matwin and C. Sammut, editors, *Proceedings of the Twelfth International Conference on Inductive Logic Programming (ILP-02)*, volume 2583 of *LNCS*, pages 133–148, Sydney, Australia, 2002. Springer.
 - [63] J. W. Lloyd. *Foundations of Logic Programming*. Springer, Berlin, 2. edition, 1989.
 - [64] C. H. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. The MIT Press, 1999.
 - [65] G. J. McKachlan and T. Krishnan. *The EM Algorithm and Extensions*. John Eiley & Sons, Inc., 1997.
 - [66] S. Muggleton. Stochastic logic programs. In L. De Raedt, editor, *Advances in Inductive Logic Programming*. IOS Press, 1996.
 - [67] S. Muggleton. Learning stochastic logic programs. *Electronic Transactions in Artificial Intelligence*, 4(041), 2000.
 - [68] S. Muggleton. Learning structure and parameters of stochastic logic programs. In S. Matwin and C. Sammut, editors, *Proceedings of the Twelfth International Conference on Inductive Logic Programming (ILP-02)*, volume 2583 of *LNCS*, pages 198–206, Sydney, Australia, 2002. Springer.
 - [69] S. Muggleton and L. De Raedt. Inductive logic programming: Theory and methods. *Journal of Logic Programming*, 19(20):629–679, 1994.
 - [70] S. H. Muggleton. Learning stochastic logic programs. In L. Getoor and D. Jensen, editors, *Working Notes of the AAAI-2000 Workshop on Learning Statistical Models from Relational Data (SRL-00)*, pages 36 –41, Austin, Texas, 2000. AAAI Press.
 - [71] R. Ng and V. S. Subrahmanian. Probabilistic Logic Programming. *Information and Computation*, 101(2):150–201, 1992.
 - [72] L. Ngo and P. Haddawy. A Knowledge-Based Model Construction Approach to Medical Decision Making. In *Proceedings of the Twentieth American Medical Informatics Association Annual Fall Symposium (AMIA-96)*, Washington DC, USA, 1996.
 - [73] L. Ngo and P. Haddawy. Answering queries from context-sensitive probabilistic knowledge bases. *Theoretical Computer Science*, 171:147–177, 1997.
 - [74] S.-H. Nienhuys-Cheng and R. de Wolf. *Foundations of Inductive Logic Programming*. Springer-Verlag, 1997.

- [75] N. J. Nilsson. *Principles of Artificial Intelligence*. Springer-Verlag, New York, 1986. Reprint, Originally published: Tioga Publishing Co. , 1980.
- [76] H. Pasula, B. Marthi, B. Milch, S. Russell, and I. Shpitser. Identity Uncertainty and Citation Matching. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*. MIT Press, 2003.
- [77] H. Pasula and S. Russell. Approximate inference for first-order probabilistic languages. In B. Nebel, editor, *Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 741–748, Seattle, Washington, USA, 2001. Morgan Kaufmann.
- [78] J. Pearl. *Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 2. edition, 1991.
- [79] A. Pfeffer and D. Koller. Semantics and Inference for Recursive Probability Models. In H. Kautz and B. Porter, editors, *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-00)*, pages 538–544., Austin, Texas, USA, 2000. AAAI Press.
- [80] A. Pfeffer, D. Koller, B. Milch, and K. T. Takusagawa. Spook: A system for probabilistic object-oriented knowledge representation. In *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-1999)*, 1999.
- [81] A. J. Pfeffer. *Probabilistic Reasoning for Complex Systems*. PhD thesis, Stanford University, 2000.
- [82] D. Poole. Probabilistic Horn abduction and Bayesian networks. *Artificial Intelligence*, 64:81–129, 1993.
- [83] F. Provost and P. Domingos. Tree induction for probability-based ranking. *Machine Learning*, 2003. (to appear).
- [84] L. R. Rabiner. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [85] L. R. Rabiner and B. H. Juang. An introduction to hidden Markov models. *IEEE ASSP Magazine*, pages 4–15, January 1986.
- [86] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Inc., 1995.
- [87] V. Santos Costa, D. Page, M. Qazi, and J. Cussens. CLP(BN): Constraint Logic Programming for Probabilistic Knowledge. In *Proceedings of the Nineteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, Mexico, 2003. Morgan Kaufman. (to appear).
- [88] T. Sato. A Statistical Learning Method for Logic Programs with Distribution Semantics. In L. Sterling, editor, *Proceedings of the Twelfth International Conference on Logic Programming (ICLP-1995)*, pages 715 – 729, Tokyo, Japan, 1995. MIT Press.
- [89] T. Sato. Parameterized logic programs where computing meets learning. In H. Kuchen and K. Ueda, editors, *Proceedings of Fifth International Symposium on Functional and Logic Programming (FLOPS-01)*, volume 2024 of *LNCS*, pages 40–60, Tokyo, Japan, 2001. Springer-Verlag.
- [90] T. Sato and Y. Kameya. PRISM: A Symbolic–Statistical Modeling Language. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, pages 1330–1339, Nagoya, Japan, 1997. Morgan Kaufmann.
- [91] T. Sato and Y. Kameya. A Viterbi-like algorithm and EM learning for statistical abduction. In R. Dybowski, J. Myers, and S. Parsons, editors, *Workshop Notes of UAI-00 Workshop on Fusion of Domain Knowledge with Data for Decision Support*, Stanford, CA, USA, 2000.
- [92] T. Sato and Y. Kameya. Parameter learning of logic programs for symbolic-statistical modeling. *Journal of Artificial Intelligence Research*, 15:391–454, 2001.
- [93] E. Segal, Y. Barash, I. Simon, N. Friedman, and D. Koller. From Promoter Sequence to Expression: A Probabilistic Framework. In *Proceedings of the Sixth International Conference on Research in Computational Molecular Biology (RECOMB-02)*, pages 263–272, Washington, DC, USA, 2002.
- [94] E. Segal, A. Battle, and D. Koller. Decomposing Gene Expression into Cellular Processes. In R. B. Altman, A. K. Dunker, L. Hunter, T. A. Jung, and T. E. Klein, editors, *Proceedings of the Pacific Symposium on Biocomputing*, pages 89 – 100, Kauai, Hawaii, USA, 2003. World Scientific.
- [95] E. Segal, B. Taskar, A. Gasch, N. Friedman, and D. Koller. Rich Probabilistic Models for Gene Expression. *Bioinformatics*, 17:S243–S252, 2001.
- [96] E. Y. Shapiro. *Algorithmic Program Debugging*. MIT Press, 1983.
- [97] E.Y. Shapiro. Logic Programs with Uncertainties: A Tool for Implementing Expert Systems. In A. Bundy, editor, *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI-1983)*, pages 529–532, Karlsruhe, Germany, 1983. William Kaufmann.
- [98] B. Taskar, E. Segal, and D. Koller. Probabilistic clustering in relational data. In B. Nebel, editor, *Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 870–87, Seattle, Washington, USA, 2001. Morgan Kaufmann.
- [99] R. J. Williams and D. Zipser. Gradient-Based Learning Algorithms for Recurrent Networks and Their Computational Complexity. In Y. Chauvin and D. E. Rumelhart, editors, *Back-propagation: Theory, Architectures and Applications*. Lawrence Erlbaum, Hillsdale, NJ: Erlbaum, 1995.
- [100] S. Wrobel. First Order Theory Refinement. In L. De Raedt, editor, *Advances in Inductive Logic Programming*. IOS Press, 1996.