

Languages for Probabilistic Modeling Over Structured and Relational Domains



Fabio Gagliardi Cozman

1 **Abstract** In this chapter we survey languages that specify probability distributions
2 using graphs, predicates, quantifiers, fixed-point operators, recursion, and other log-
3 ical and programming constructs. Many of these languages have roots both in prob-
4 abilistic logic and in the desire to enhance Bayesian networks and Markov random
5 fields. We examine their origins and comment on various proposals up to recent
6 developments in probabilistic programming.

7 1 Introduction

8 Diversity is a mark of research in artificial intelligence (AI). From its inception, the
9 field has exercised freedom in pursuing formalisms to handle monotonic, nonmono-
10 tonic, uncertain, and fuzzy inferences. Some of these formalisms have seen cycles
11 of approval and disapproval; for instance, probability theory was taken, in 1969, to
12 be “epistemologically inadequate” by leading figures in AI (McCarthy and Hayes
13 1969). At that time there was skepticism about combinations of probability and logic,
14 even though such a combination had been under study for more than a century.

15 A turning point in the debate on the adequacy of probability theory to AI was
16 Judea Pearl’s development of Bayesian networks (Pearl 1988). From there many
17 other models surfaced, based on the notion that we must be able to specify probab-
18 ility distributions in a modular fashion through independence relations (Sadeghi and
19 Lauritzen 2014). In spite of their flexibility, Bayesian networks are “propositional”
20 in the sense that random variables are not parameterized, and one cannot quantify
21 over them. For instance, if you have 1000 individuals in a city, and for each of them
22 you are interested in three random variables (say education, income, and age), then
23 you must explicitly specify 3000 random variables and their independence relations.

24 There have been many efforts to extend graphical models so as to allow one to
25 encode repetitive patterns, using logical variables, quantification, recursion, loops,
26 and the like. There are specification languages based on database schema, on first-

F. G. Cozman (✉)
University of São Paulo, São Paulo, Brazil
e-mail: fgcozman@usp.br

© Springer Nature Switzerland AG 2020
P. Marquis et al. (eds.), *A Guided Tour of Artificial Intelligence Research*,
https://doi.org/10.1007/978-3-030-06167-8_9

1

27 order logic, on logic programming, on functional programming, even on procedural
 28 programming. Often these languages employ techniques from seminal probabilistic
 29 logics. The purpose of this chapter is to review some of these languages, starting with
 30 a brief review of probabilistic logic concepts, and then moving to relational variants
 31 of Bayesian networks and to probabilistic programming.

32 In Sect. 2 we present some foundational results from probabilistic logic, so as to
 33 fix useful terminology concerning syntax and semantics. In Sect. 3 we look at sev-
 34 eral formalisms that mix, using graphs, Bayesian networks and relational modeling.
 35 Section 4 is devoted to probabilistic logic programming; in Sect. 5 we go through lan-
 36 guages inspired by various logics, and in Sect. 6 we examine Markov random fields.
 37 In Sect. 7 we consider the emerging field of probabilistic programming. Finally, in
 38 Sect. 8 we very briefly mention some inference and learning techniques.

39 2 Probabilistic Logics: The Laws of Thought?

40 Boole’s book on The Laws of Thought is aptly sub-titled “on which are founded
 41 the mathematical theories of logic and probabilities” (Boole 1958). Starting from
 42 that effort, many other thinkers examined the combination of first-order logic and
 43 probabilities (Gaifman 1964; Gaifman and Snir 1982; Hoover 1978; Keisler 1985;
 44 Scott and Krauss 1966). A mix of logical and probabilistic reasoning was also central
 45 to de Finetti’s concept of probability (Coletti and Scozzafava 2002; de Finetti 1964).
 46 Nilsson (1986) later rediscovered some of these ideas in the context of AI, in particular
 47 emphasizing linear programming methods that had been touched before (Bruno and
 48 Gilio 1980; Hailperin 1976).

49 At the beginning of the nineties, Nilsson’s probabilistic logic and its extensions
 50 were rather popular amongst AI researchers (Hansen and Jaumard 1996); in particu-
 51 lar, probabilistic first-order logic received sustained attention (Bacchus 1990; Fagin
 52 et al. 1990; Halpern 2003). That feverish work perhaps convinced some that the laws
 53 of thought had indeed been nailed down.

54 We now review some concepts used in probabilistic logics, as they are relevant to
 55 the remainder of this survey.

56 Propositional logic consists of formulas containing *propositions* A_1, A_2, \dots , and
 57 the *Boolean operators* \wedge (conjunction), \vee (disjunction), \neg (negation), \rightarrow (implica-
 58 tion) and \leftrightarrow (equivalence); further details are discussed in chapter “Reasoning with
 59 Propositional Logic - From SAT Solvers to Knowledge Compilation” of this Vol-
 60 ume. First-order logic consists of formulas containing predicates and functions, plus
 61 the Boolean operators, logical variables and existential/universal quantifiers (further
 62 details are discussed in chapter “Automated Deduction” of this Volume). Any pred-
 63 icate r , and any function f , is associated with a nonnegative integer, its *arity*. A
 64 predicate of arity zero is treated as a proposition. A function of arity zero is a *con-*
 65 *stant*. A *term* is either a logical variable or a constant. A predicate of arity k , followed
 66 by k terms (usually in parenthesis) is an *atom*. For instance, if *baker* is a predicate
 67 of arity 1, then both *baker*(λ) and *baker*(John) are atoms.

68 The syntax of a minimal propositional probabilistic logic is rather simple: we
 69 can have any propositional formula ϕ , and moreover any *probabilistic assessment*
 70 $\mathbb{P}(\phi) = \alpha$, where ϕ is a propositional formula and α is a number in $[0, 1]$. For
 71 instance, both $A_1 \wedge \neg A_2$ and $\mathbb{P}(\neg A_3 \vee A_4 \vee \neg A_5) = 1/2$ are well-formed formulas.
 72 Conditional probabilistic assessments are often allowed; that is, $\mathbb{P}(\phi|\theta) = \alpha$ where
 73 ϕ and θ are propositional formulas.

74 *Example 1* Suppose A_1 means “Tweety is a penguin”, A_2 means “Tweety is a bird”,
 75 and A_3 means “Tweety flies”. Then

$$76 \quad A_1 \rightarrow A_2, \quad A_1 \rightarrow \neg A_3, \quad \mathbb{P}(A_3|A_2) = 0.95$$

77 is a set of well-formed formulas. □

78 The usual semantics associated with this propositional syntax is as follows. Sup-
 79 pose we have propositions A_1, A_2, \dots, A_n . There are 2^n *truth assignments* to these
 80 propositions (each proposition assigned **true** or **false**). To simplify the terminology,
 81 we refer to a truth assignment as an *interpretation*. Propositional probabilistic logic
 82 focuses on probability measures over interpretations, where the sample space Ω is
 83 the set of 2^n interpretations. Recall that such a measure \mathbb{P} can be specified by associ-
 84 ating each element of Ω with a nonnegative number in $[0, 1]$, guaranteeing that these
 85 numbers add up to one (as discussed in chapter “Representations of Uncertainty in
 86 Artificial Intelligence: Probability and Possibility” of Volume 1).

87 A probabilistic assessment $\mathbb{P}(\phi) = \alpha$, for some $\alpha \in [0, 1]$ is interpreted as a
 88 constraint; namely, the constraint that the probability of the set of interpretations
 89 satisfying ϕ is exactly α . A conditional probability assessment $\mathbb{P}(A_1|A_2) = \alpha$
 90 is usually read as the constraint $\mathbb{P}(A_1 \wedge A_2) = \alpha \mathbb{P}(A_2)$.

91 *Example 2* Consider a (rather simplified) version of Boole’s challenge problem
 92 (Hailperin 1996): we have $\mathbb{P}(A_1) = \alpha_1$ and $\mathbb{P}(A_2) = \alpha_2$, and moreover $\mathbb{P}(A_3|A_1) =$
 93 β_1 and $\mathbb{P}(A_3|A_2) = \beta_2$; finally we know that $A_3 \rightarrow (A_1 \vee A_2)$. What are the possible
 94 values for $\mathbb{P}(A_3)$?

95 There are three propositions A_1, A_2, A_3 ; an interpretation can be encoded as a
 96 triple $(a_1 a_2 a_3)$ where $a_i = 1$ when A_i is **true**, and $a_i = 0$ otherwise. There are 8
 97 interpretations that might be possible, but interpretation (001) is impossible because
 98 $A_3 \rightarrow (A_1 \vee A_2)$ holds. Each interpretation is to have a probability; we denote by p_j
 99 the probability of the j th interpretation, where we order lexicographically the triple
 100 $(a_1 a_2 a_3)$ as if it were a binary number. Thus the constraint $\mathbb{P}(A_1) = \alpha_1$ means

$$101 \quad p_4 + p_5 + p_6 + p_7 = \alpha_1,$$

102 and the constraint $\mathbb{P}(A_3|A_2) = \beta_2$ means

$$103 \quad p_3 + p_7 = \beta_2(p_2 + p_3 + p_6 + p_7).$$

104 By minimizing/maximizing $\mathbb{P}(A_3) = p_3 + p_5 + p_7$, subject to these constraints and
 105 $p_j \geq 0$ for all j and $\sum_j p_j = 1$, we obtain $\mathbb{P}(A_3) \in [L, U]$, where

$$L = \max(\alpha_1\beta_1, \alpha_2\beta_2),$$

$$U = \min(\alpha_1\beta_1 + \alpha_2\beta_2, \alpha_1\beta_1 + (1 - \alpha_1), \alpha_2\beta_2 + (1 - \alpha_2)),$$

provided $0 \leq L \leq U \leq 1$ (otherwise the whole problem is inconsistent). \square

To use de Finetti’s terminology, a set of formulas and assessments that is satisfied by at least one probability measure is said to be *coherent* (Coletti and Scozzafava 2002). Given a coherent set of assessments, reasoning should only inform us about the least commitment conclusions that are necessarily correct.

We can of course contemplate probabilistic assessments $\mathbb{P}(\phi) = \alpha$ where ϕ is a formula of first-order logic and α is a number in $[0, 1]$.¹ The semantics of first-order formulas is given by a *domain* and an *interpretation*. A domain \mathcal{D} , in this technical sense, is just a set (and should not be taken as the sort of “structured domain knowledge” alluded to in the title of this chapter). An interpretation is a mapping from each predicate of arity k to a relation in \mathcal{D}^k , and from each function of arity k to a function from \mathcal{D}^k to \mathcal{D} (Enderton 1972).

Each probabilistic assessment $\mathbb{P}(\phi) = \alpha$, where ϕ is a first-order formula, is interpreted as a constraint on the probability measures over the set of interpretations for a fixed domain. If the domain is infinite, the set of interpretations is uncountable. In this overview we can bypass difficulties with uncountable sets, but still present the main points, by assuming all domains to be finite, and moreover by assuming that no functions, other than constants, are present.

Example 3 Consider predicates *penguin*, *bird*, and *flies*, all of arity 1, and predicate *friends*, of arity 2. The intended meaning is that *penguin*(Tweety) indicates that Tweety is a penguin, and likewise for *bird*, while *flies*(Tweety) indicates that Tweety flies, and *friends*(Tweety, Skipper) indicates that Tweety and Skipper are friends. Suppose the domain is $\mathcal{D} = \{\text{Tweety, Skipper, Tux}\}$. An interpretation might assign both Tweety and Tux to *penguin*, only the pair (Tweety, Skipper) to *friends*, and so on.

Suppose $\mathbb{P}(\forall\chi : \exists y : \text{friends}(\chi, y)) = 0.01$. This assigns probability 0.01 to the set of interpretations where any element of the domain has a friend. Another possible assessment is $\forall\chi : \mathbb{P}(\text{penguin}(\chi)) = 0.03$; note that here the quantifier is “outside” of the probability.

For a domain with N elements, we have 2^N possible interpretations for *penguin*, and 2^{N^2} interpretations for *friends*; the total number of possible interpretations for the predicates is 2^{3N+N^2} . \square

We refer to the semantics just outlined as an *interpretation-based* semantics, because probabilities are assigned to sets of interpretations. There is also a *domain-based* semantics, where we assume that probability measures are defined over the domain. This may be useful to capture some common scenarios. For instance, consider: “The probability that some citizen is a computer scientist is α ”. We might wish

¹We might be more even general by introducing “probabilistic quantifiers”, say by writing $\mathbb{P}^{\geq\alpha}\phi$ to mean $\mathbb{P}(\phi) \geq \alpha$. We could then nest $\mathbb{P}^{\geq\alpha}\phi$ within other formulas (Halpern 2003). We avoid this generality here.

144 to interpret this through a probability measure over the set of citizens (the domain);
 145 the constraint is that the set of computer scientists gets probability α . Domain-based
 146 semantics, and even combinations of interpretation- and domain-based semantics,
 147 have been investigated for a while (Bacchus 1990; Fagin et al. 1990; Hoover 1978;
 148 Keisler 1985); however, interpretation-based semantics are more commonly adopted.

149 First-order probabilistic logic has high representational power, but very high complex-
 150 ity (Abadi and Halpern 1994). Another drawback is *inferential vacuity*: given a
 151 set of formulas and assessments, usually all that can be said about the probability
 152 of some other formula is that it lies in a large interval, say between 0.01 and 0.99.
 153 This happens because there are exponentially many interpretations, and a few assess-
 154 ments do not impose enough constraints on probability values. Finally, there is yet
 155 another problem: it has been observed that first-order logic itself is not sufficient to express
 156 recursive concepts or default assumptions, and tools for knowledge representation
 157 often resort to special constructs (Baader and Nutt 2002; Baral 2003). Hence it does
 158 not seem that the machinery of first-order probabilistic logic, however elegant, can
 159 exhaust the laws of thought, after all.

160 3 Bayesian Networks and Their Diagrammatic 161 Relational Extensions

162 Bayesian networks offer a pleasant way to visualize independence relations, and an
 163 efficient way to encode a probability distribution; as such, they have been widely
 164 applied within AI and in many other fields (Darwiche 2009; Koller and Friedman
 165 2009; Pourret et al. 2008). To fix terminology, here is a definition (see also chapter
 166 “Belief Graphical Models for Uncertainty Representation and Reasoning” of this
 167 Volume). A Bayesian network consists of a pair: there is a directed acyclic graph
 168 \mathbb{G} , where each node is a random variable X_i , and a probability distribution \mathbb{P} over
 169 the random variables, so that \mathbb{P} satisfies the Markov condition with respect to \mathbb{G} :
 170 each X_i is independent of its nondescendants (in \mathbb{G}) given its parents (in \mathbb{G}) (Pearl
 171 1988). Even though we can have discrete and continuous random variables in a
 172 Bayesian network, in this survey we simplify matters by focusing on *binary* random
 173 variables. When we have a finite set of binary random variables, the Markov condition
 174 implies a factorization of the joint probability distribution; for any configuration
 175 $\{X_1 = x_1, \dots, X_n = x_n\}$,

$$176 \quad \mathbb{P}(X_1 = x_1, \dots, X_n = x_n) = \prod_{i=1}^n \mathbb{P}(X_i = x_i | \text{pa}(X_i) = \pi_i),$$

177 where $\text{pa}(X_i)$ denotes the parents of X_i , π_i is the projection of $\{X_1 = x_1, \dots, X_n =$
 178 $x_n\}$ on $\text{pa}(X_i)$. Often each $\mathbb{P}(X_i = x_i | \text{pa}(X_i) = \pi_i)$ is described by a *local condi-*
 179 *tional probability table*.




Fig. 1 The Bayesian network for the “propositional” version of the University World, where $p(1|a, b)$ denotes $\mathbb{P}(\text{JohnFailsCourse} = 1 | \text{JohnIsDedicated} = a, \text{CourseIsHard} = b)$

180 Consider, as an example, a simplified version of the ubiquitous “University World”
 181 (Getoor et al. 2007). We have random variables **JohnIsDedicated**, **CourseIsHard**,
 182 and **JohnFailsCourse**, each one with values 0 and 1. Figure 1 depicts a Bayesian
 183 network, where **JohnIsDedicated** and **CourseIsHard** are independent, and where
 184 both directly affect **JohnFailsCourse**.

185 Bayesian networks do encode structured domain knowledge through its independent
 186 relations. However, domain knowledge may come with much more structured
 187 patterns. For example, a University World usually has many students and many
 188 courses, and a very repetitive structure. We might say: for *any* pair (χ, y) , where χ is
 189 a student and y is a course, the probability that χ fails given that she is dedicated and
 190 the course is easy is 0.1. Figure 2 depicts a Bayesian network with various students
 191 and courses, where probabilities are obtained by repeating the conditional probability
 192 tables in Fig. 1.

193 The question then arises as to how we should specify such “very structured”
 194 Bayesian networks. It makes sense to import some tools from first-order logic. For
 195 instance, we clearly have predicates, such as the predicate **fails**, that can be *grounded*
 196 by replacing logical variables by elements of appropriate domains (thus we obtain the
 197 grounded predicate **fails(Tina, Physics)**, and so on). However, here the “grounded
 198 predicates” that appear in a graph are not just propositions, but rather random vari-
 199 ables. A symbol such as **fails** must be understood with a dual purpose: it can be
 200 viewed as a predicate, or as a function that yields a random variable for each pair
 201 of elements of the domain. And a symbol such as **fails(Tina, Physics)** also has a
 202 dual purpose: it can be viewed as a grounded predicate, or as a random variable that
 203 yields 1 when the grounded predicate is **true**, and 0 otherwise.

204 We adopt the following terminology (Poole 2003). A *parvariable* (for *parameter-*
 205 *ized random variable*) is a function that yields a random variable (its *grounding*) when
 206 its parameters are substituted for elements of the domain. The number of parameters
 207 of a parvariable is its *arity*. To specify a parameterized Bayesian network, we use
 208 parvariables instead of random variables.

209 Of course, when we have parvariables we must adapt our conditional probability
 210 tables accordingly, as they depend on the values of parvariables. For example, for
 211 the Bayesian network in Fig. 2 we might have

$$\mathbb{P}(\text{isDedicated}(\chi) = 1) = 0.6,$$

212

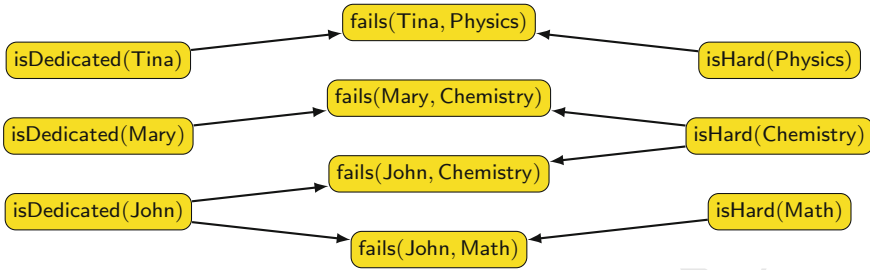


Fig. 2 The Bayesian network for the University World with three students and three courses

213 meaning $\mathbb{P}(\text{isDedicated}(a)) = 0.6$ for each student a . Also, we might write

214
$$\mathbb{P}(\text{fails}(\chi, y) = 1 | \text{isDedicated}(\chi) = 0, \text{isHard}(y) = 0) = 0.5,$$

215
$$\mathbb{P}(\text{fails}(\chi, y) = 1 | \text{isDedicated}(\chi) = 0, \text{isHard}(y) = 1) = 0.8,$$

 216

217 and so on, assessments that are imposed on every pair (χ, y) .

218 Each “parameterized table” specifying a conditional probability table for each
 219 substitution of logical variables is called a *parfactor*. Thus for the Bayesian network
 220 in Fig. 2 we need only three parfactors.

221 Possibly the most popular diagrammatic scheme to specify parvariables and par-
 222 factors is offered by *plate models*. A plate consists of a set of parvariables that
 223 share a domain (that is, the parvariables are all indexed by elements of a domain).
 224 A plate model for the University World is presented in Fig. 3 (left); plates are usu-
 225 ally depicted as rectangles enclosing the related parvariables. The main constraint
 226 imposed on plates is that the domains used to index the parents of a parvariable must
 227 also be used to index the parvariable. Thus in Fig. 3 (left) we see that *fails* appears
 228 in the intersection of two plates, each one of them associated with a domain: *fails* is
 229 indexed by both domains.

230 Plate models appeared within the BUGS project (Gilks et al. 1993; Lunn et al.
 231 2009) around 1992. At that time other template languages were discussed under the
 232 general banner of “knowledge-based model construction” (Goldman and Charniak
 233 1990; Horsch and Poole 1990; Wellman et al. 1992). Plates were promptly adopted
 234 in machine learning (Buntine 1994) and in many statistical applications (Lunn et al.
 235 2012).

236 The BUGS package was innovative not only in proposing plates but also in intro-
 237 ducing a textual language in which to specify large Bayesian networks through plates.
 238 Figure 3 (right) shows a plate model rendered by WinBUGS (a version of BUGS);
 239 this plate model can be specified textually by using a loop to describe the plate, and
 240 by introducing a statement for each parvariable in the plate model:

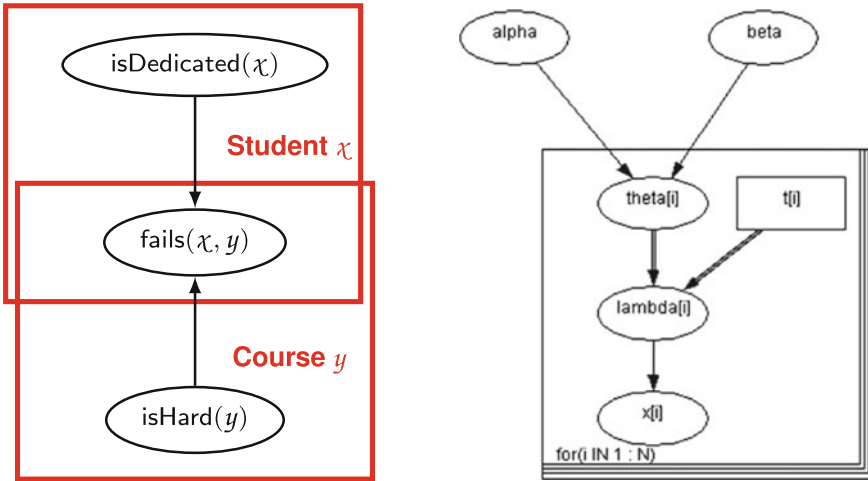


Fig. 3 Left: Plate model for the University World. Right: A plate model rendered in WinBUGS (described in the WinBUGS manual)

```

241  model {
242    for (i in 1 : N) {
243      theta[i] ~ dgamma(alpha, beta)
244      lambda[i] <- theta[i] * t[i]
245      x[i] ~ dpois(lambda[i])
246    }
247    alpha ~ dexp(1)
248    beta ~ dgamma(0.1, 1.0)
249  }

```

250 The symbol \sim indicates that the left hand side is distributed according to the distribu-
 251 tion in the right hand side (Gamma, Poisson, Exponential distributions, respectively
 252 specified by `dgamma`, `dpois`, `dexp`), while the symbol `<-` indicates a deterministic
 253 expression. Rectangular nodes in the plate model, such as the one containing `t[i]`,
 254 denote constants that are specified elsewhere in the program. BUGS is particularly
 255 powerful in that it can go beyond finite modeling, allowing discrete and continuous
 256 distributions (Lunn et al. 2012).

257 A possible extension of plate models is to allow a parvariable to have children
 258 outside of its plate. Figure 4 presents a popular example of such an “enhanced”
 259 plate; many existing topic models and factorization schemes are similarly drawn
 260 with “enhanced” plates.

261 When a parvariable does not belong to the plate of its parent parvariables, one must
 262 worry about *aggregation*. To understand this, consider an example. Suppose that in the
 263 University World we have a parvariable `highFailureRate`, of arity zero, whose value
 264 depends on `fails(χ, y)` for all pairs (χ, y) . This parvariable should be drawn outside




Fig. 4 The usual “enhanced” plate model for *smoothed Latent Dirichlet Allocation (sLDA)* (Blei et al. 2003). Logical variables and domains are omitted. Here X is the parent of W , but X and W belong to non-intersecting plates

265 of all plates in Fig. 3 (left). Consider specifying the parfactor for `highFailureRate`.
 266 All the groundings of `fails(χ , y)`, for all pairs (χ, y) , affect `highFailureRate`. Thus
 267 to specify probabilities for the latter we must somehow *aggregate* values of the
 268 former. This is akin to quantification in first-order logic. Most languages that allow
 269 probability values to depend on several objects at once resort to so-called *combination*
 270 *functions* (some of them will be discussed later).

271 Plates were not the only tools proposed in the nineties to specify repetitive
 272 Bayesian networks; *network fragments* and *object-oriented networks* were also con-
 273 templated (Glesner and Koller 1995; Koller and Pfeffer 1997, 1998; Mahoney and
 274 Laskey 1996). These ideas evolved to *Probabilistic Relational Models (PRMs)*, a
 275 clever mix of Bayesian networks and entity-relationship diagrams that spurred many
 276 efforts on knowledge representation and on *Statistical Relational Learning* (Fried-
 277 man et al. 1999; Getoor et al. 2007; Koller and Pfeffer 1998).

278 In short, a PRM consists of a set of classes; each class contains a set of objects
 279 (similarly to a domain), and is associated with a set of parvariables. Usually a class
 280 is drawn as a box containing parvariables. Figure 5 (top) depicts a possible PRM for
 281 the University World; each class contains a parvariable, and there are *association*
 282 edges between classes, indicating for instance that each **Registration** is paired with
 283 a **Student** and with a **Course**. Associations appear as dashed edges; in PRMs these
 284 associations are often called *slot-chains* (Getoor et al. 2007).

285 Each parvariable X in a PRM is then associated with a parfactor specifying the
 286 probability for each instance of X given instances of parents of X . For instance,
 287 for the University World PRM we would need a parfactor to specify probability
 288 of `fails(z)` as dependent on the value of appropriate instances of `isDedicated(χ)`
 289 and `isHard(y)`. Here “appropriate instance” means “instance related by appropriate
 290 association”. It may happen that a parvariable depends on many instances of another
 291 parvariable; for example, we may have a parvariable `highFailureRate` that depends
 292 on all groundings of `fails`. In PRMs a parfactor may depend on an *aggregate* of
 293 instances (an aggregate is just a combination function).

294 A *relational skeleton* for a PRM is an explicit specification of objects in each class,
 295 plus the explicit specification of pairs of objects that are associated. The semantics
 296 of a PRM with respect to a skeleton is the Bayesian network that is produced by
 297 instantiating the parvariables and parfactors.

298 Suppose we build a graph where each node is a parvariable, and the parents of
 299 a parvariable r are the parvariables that appear in the parfactor for r . If this graph
 300 is acyclic, then every skeleton will induce an acyclic (thus valid) Bayesian network.

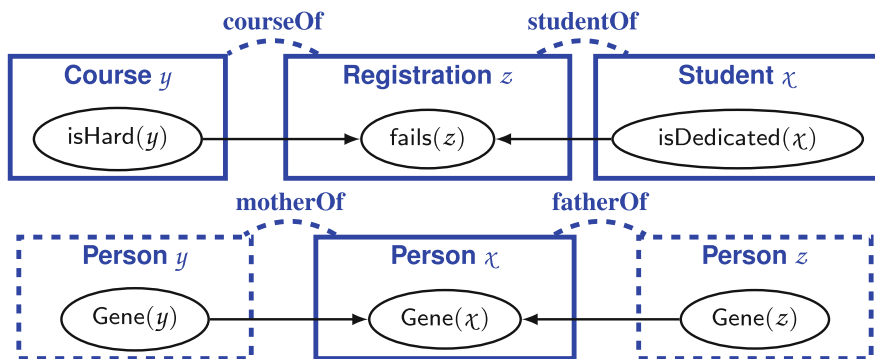


Fig. 5 Top: PRM for the University World. Bottom: PRM representing genetic relationships

301 However, if that graph has cycles, we face a *global semantics* problem (Jaeger 2002):
 302 can we guarantee that an acyclic Bayesian network emerges for any skeleton we
 303 expect to have? Suppose for example that a particular gene in a person may depend on
 304 that gene in the person’s father and mother. Some specification languages for PRMs
 305 allow classes to appear more than once in a diagram, as depicted in Fig. 5 (bottom).
 306 The dependences may look cyclic, but we know that “fatherOf” and “motherOf”
 307 are acyclic relations, so no cycles can appear when we consider specific skeletons.
 308 Algorithms that decide existence of global semantics have been provided for special
 309 cases (Getoor et al. 2007; Milch et al. 2005b), but the answer to this question in full
 310 generality seems quite hard, and possibly undecidable (De Bona and Cozman 2017;
 311 Jaeger 2002).

312 As a digression, note that we can think of associations themselves as random
 313 variables that are fully observed in the relational skeleton. For instance, we might
 314 have a parvariable *isCourseOf* that indicates whether a course y is actually in a
 315 particular registration z . By grounding such a specification, one can produce a ground
 316 Bayesian network as depicted in Fig. 2, where only a few selected pairs student/course
 317 are associated with a grade. A parvariable that corresponds to an association, and
 318 that is fully observed in the relational skeleton, is called a *guard parvariable* (Koller
 319 and Friedman 2009).

320 There is much in favor of PRMs: they are as modular and visually elegant as
 321 Bayesian networks, and as featured as entity-relationship diagrams. They allow one
 322 to represent uncertainty about the structure, by associating probabilities with associa-
 323 tions; they even allow for uncertainty about existence of particular instances (Getoor
 324 et al. 2007). Moreover, inference and learning algorithms for Bayesian networks can
 325 be easily adapted to PRMs (Sect. 8).

326 However, a drawback of PRMs is that they have no unified formal syntax; indeed
 327 the term “PRM” has a rather loose meaning and it is difficult to draw boundaries
 328 on what is, and what is not, a PRM. There have been some attempts to formalize
 329 PRMs. One of them is the DAPER language, where entity-relationship diagrams are
 330 extended to cope with probabilities (Heckerman et al. 2007). Figure 6 shows DAPER
 331 diagrams for the University World, for the genetic problem, and for a Hidden Markov

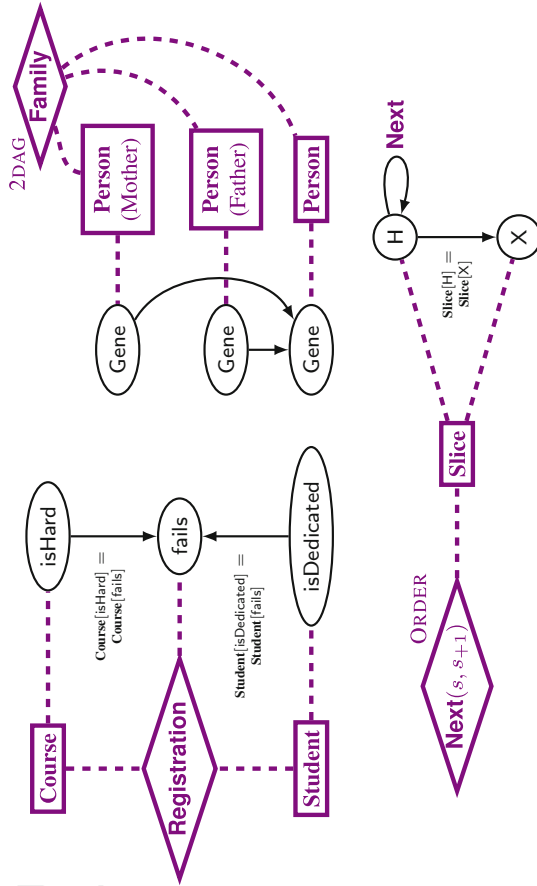


Fig. 6 Top left: A DAPER diagram for the University World. Top right: A DAPER diagram for the genetic problem. Bottom: A simplified DAPER diagram for a Hidden Markov Model (H denotes the state, while X denotes the observation; at each time slice there is a state transition that depends on the previous state)

332 Model, a rather simple Dynamic Bayesian network (Heckerman et al. 2007). Note that
 333 temporal modeling is in essence a relational problem, where time steps index random
 334 variables; we later consider other formalisms with a focus on temporal evolution. A
 335 DAPER diagram can be annotated with a variety of constraints, such as the 2DAG in
 336 Fig. 6 (meaning that each child of a node has at most two parents and cannot be its
 337 own ancestor!). As edges can be annotated with constraints in first-order logic, there
 338 are few guarantees one can offer with respect to the complexity of manipulating a
 339 DAPER diagram.

340 There are also textual languages that aim at encoding PRMs in a formal manner;
 341 for instance, the Probabilistic Relational Language (PRL) (Getoor and Grant 2006)
 342 and also the CLP(\mathcal{BN}) language (Costa et al. 2003) resort to logic programming to
 343 specify parvariables and parfactors.

344 Another comprehensive framework that combines logical constructs and Bayesian
 345 networks is conveyed by *Multi-Entity Bayesian Networks (MEBNs)* (da Costa and
 346 Laskey 2005; Laskey 2008). An MEBN consists of a set of network fragments (called
 347 *MFrag*s), each containing parvariables (referred to as *template random variables*)
 348 and constraints over logical variables, all with associated parfactors. The language
 349 is very expressive as it allows for parfactors that are specified programmatically, and
 350 constraints that are based on first-order formulas. But MEBNs go beyond first-order
 351 logic in several aspects: instead of just **true** and **false**, in MEBNs some variables may
 352 assume the value **absurd**; moreover, there are ways to specify default probabilities.
 353 The modeling framework has been implemented and applied to practical problems
 354 (Carvalho et al. 2010). Figure 7 shows an MEBN for the University World: each
 355 MFrag corresponds to a parvariable in the original plate model. Each MFrag also
 356 contains *context parvariables*, depicted in Fig. 7 in rectangles. These correspond to
 357 guard parvariables that induce the structure of any instantiated Bayesian network.
 358 Finally, each MFrag contains *input parvariables*, whose parfactors are specified in
 359 other MFrag; input variables appear in Fig. 7 as dark ovals.

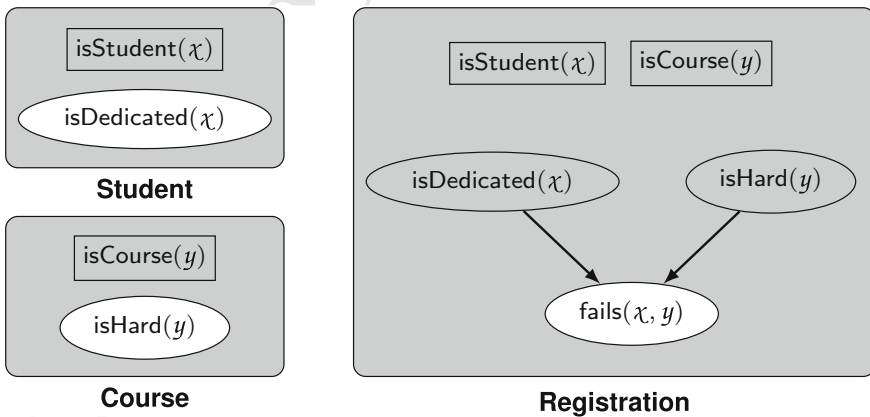


Fig. 7 An MEBN, with three network fragments, for the University World

360 The formalisms reviewed in this section rely on graphs to encode classes, associa-
 361 tions, parvariables and their dependences; they do not offer much syntactic guidance
 362 as to how one should encode the parfactors. Most of the languages we discuss in the
 363 remainder of this chapter can specify both the dependences amongst parvariables
 364 and the parfactors; to do so, they use textual descriptions.

365 4 Probabilistic Logic Programming

366 Logic programming is an old and core AI technology (Baral 2003), as discussed in
 367 chapter “Logic Programming” of this Volume. It is not surprising that probabilistic
 368 logic programming has been pursued for some time (Lukasiewicz 1998; Lakshmanan
 369 and Sadri 1994; Ng and Subrahmanian 1992; Ngo and Haddawy 1997).

370 Poole was an early advocate of the idea that logic programming can be used
 371 to extend Bayesian networks with relational modeling (Poole 1993b). In Poole’s
 372 Probabilistic Horn Abduction (PHA) language one can write:

```
373 symptom(S) <- carries(D), causes(D,S).  
374 disjoint([causes(D,S): 0.7, nc(D,S): 0.3]).
```

375 Here we have a rule, indicating how a symptom shows up, and a statement indicating
 376 that either `causes(D, S)` is true with probability 0.7, or `nc(D, S)` is true with
 377 probability 0.3. The PHA language later evolved into Independent Choice Logic
 378 (ICL), where decision-making and multi-agent scenarios can be modeled (Poole
 379 1997, 2008).

380 Another seminal proposal for a mixture of logic programming and probabilities
 381 was Sato’s *distribution semantics*, embodied in the popular package PRISM (Sato
 382 1995; Sato and Kameya 2001). A similar syntax and semantics appeared in Fuhr’s
 383 work on Probabilistic Datalog (Fuhr 1995). The distribution semantics has been
 384 adopted by many languages (Riguzzi et al. 2014; De Raedt and Kimmig 2015); to
 385 understand it, a few definitions are needed.

386 A logic program is a set of declarative rules that just describe a problem; finding
 387 a solution is the task of an inference engine (Dantsin et al. 2001; Eiter et al. 2009).
 388 A *normal logic program* consists of rules written as (Dantsin et al. 2001)

389
$$A_0 :- A_1, \dots, A_m, \mathbf{not} A_{m+1}, \dots, \mathbf{not} A_n.$$

390 where the A_i are atoms and **not** indicates negation as failure. The *head* of this rule
 391 is A_0 ; the remainder of the rule is its *body*. A rule without a body, written simply
 392 as A_0 , is a *fact*. A *subgoal* in the body is either an atom A (a *positive* subgoal) or
 393 **not** A (a *negative* subgoal). A program without negation is *definite*, and a program
 394 with only grounded atoms is *propositional*.

395 The *Herbrand base* of a program is the set of all grounded atoms built from
 396 constants and predicates in the program. As before, we allow for constants, but we
 397 do not consider functions, so as to stay with finite Herbrand bases. The grounding of a
 398 program is the propositional program obtained by applying every possible grounding
 399 to each rule, using only the constants in the program (i.e., using only grounded atoms
 400 in the Herbrand base).

401 The *dependency graph* of a program is a directed graph where each predicate is a
 402 node, and where there is an edge from a node B to a node A if there is a rule where
 403 A appears in the head and B appears in the body; if B appears right after **not** in any
 404 such rule, the edge is *negative*; otherwise, it is *positive*. The *grounded dependency*
 405 *graph* is the dependency graph of the propositional program obtained by grounding.

406 A program is *acyclic* when its grounded dependency graph is acyclic. The seman-
 407 tics of an acyclic program is rather simple, and given by the program's *Clark com-*
 408 *pletion* (Clark 1978): roughly, take the grounding of the program, and for each head
 409 A , make it **true** if some rule with head A has all its subgoals recursively assigned to
 410 **true**. At the end of this process, we have an interpretation for the predicates.

411 Take a probabilistic logic program to consist of *probabilistic facts* in addition to
 412 rules and facts. A probabilistic fact is written as $\alpha::A.$, where α is a number in $[0, 1]$
 413 and A is a fact (here we use the syntax of the popular package Problog (Fierens et al.
 414 2014)).

415 Sato's distribution semantics is, in essence, a distribution over logic programs: for
 416 each probabilistic fact $\alpha::A.$, with probability α fact A is added to the program, and
 417 with probability $1 - \alpha$ the fact is not added to the program (all probabilistic facts are
 418 selected independently). Note that each selection of probabilistic facts generates a
 419 logic program.

420 *Example 4* To understand the distribution semantics, consider a probabilistic logic
 421 program in Problog's syntax. The first line establishes some deterministic facts,
 422 followed by probabilistic ones, and then we have a few rules (each rule uses an
 423 auxiliary predicate associated with a probabilistic fact):

```

424 isStudent(john). isStudent(mary). isCourse(math).
425 0.6::isDedicated(X).
426 0.4::isHard(Y).
427 0.5::a1(X,Y). 0.1::a2(X,Y).
428 0.8::a3(X,Y). 0.4::a4(X,Y).
429 fails(X,Y) :- isStudent(X), isDedicated(X),
430               isCourse(Y), isHard(Y), a1(X,Y).
431 fails(X,Y) :- isStudent(X), isDedicated(X),
432               isCourse(Y), not isHard(Y), a2(X,Y).
433 fails(X,Y) :- isStudent(X), not isDedicated(X),
434               isCourse(Y), isHard(Y), a3(X,Y).
435 fails(X,Y) :- isStudent(X), not isDedicated(X),
436               isCourse(Y), not isHard(Y), a4(X,Y).
```

437 The Herbrand base is then the set of groundings obtained by replacing logical vari-
 438 ables by the constants `john`, `mary`, and `math`. Altogether, these facts, probabilistic
 439 facts, and rules are similar to the plate mode in Fig. 3 (left). Note that predicates
 440 `isStudent` and `isCourse` correspond to guard parvariables; they can also be seen as
 441 the context parvariables depicted in Fig. 7. \square

442 One convenient feature of probabilistic logic programs (with Sato’s distribution
 443 semantics) is that they inherit the “default” assumptions used in logic programming.
 444 For instance, if we do not say that `isCourse(mary)` is true, then it is automatically
 445 false. This simplifies the specification of “guard parvariables”: for instance, in the
 446 example above the parvariables `isStudent` and `isCourse` aptly specify the structure of
 447 the graph, and are not associated with probabilities.

448 Poole’s PHA focused on acyclic probabilistic logic programs, while Sato’s original
 449 proposal focused on definite, but not necessarily acyclic, probabilistic logic programs.
 450 These syntactic restrictions have been removed in a variety of ways. One natural
 451 extension is to allow *stratified negation*; that is, to allow cycles in the dependency
 452 graph as long as there is no negative edge in any cycle (Dantsin et al. 2001). Here is
 453 a well-formed probabilistic logic program in Problog’s syntax:

```
454     smokes(X) :- not relaxed(X) .
455     smokes(X) :- influences(Y,X), smokes(Y) .
456     0.3 :: influences(john,mary) .
457     0.3 :: influences(mary,john) .
458     0.2 :: relaxed(mary) .
```

459 For any sample of the three probabilistic facts, the resulting logic program has a
 460 cycle. But this is not a problem, for cyclic logic programs still have a semantics
 461 (Dantsin et al. 2001); in fact, any probabilistic logic program such that the rules form
 462 a stratified logic program induces a *unique* probability measure over interpretations.
 463 The Problog package deals with probabilistic stratified logic programs (Fierens et al.
 464 2014; De Raedt and Kimmig 2015).

465 Probabilistic logic programs that contain cycles and arbitrary negation have been
 466 also studied (Hadjichristodoulou and Warren 2012; Lukasiewicz 2005; Riguzzi 2015;
 467 Sato et al. 2005). Several semantics have been proposed for such programs. We
 468 mention here two semantics: one is based on the *stable model* semantics, where a
 469 logic program admits more than a single stable model as its meaning (Gelfond and
 470 Lifschitz 1988), and the *well-founded semantics*, where a logic program admits a
 471 single well-founded model, but that model may contain true, false and undefined
 472 assignments (Van Gelder et al. 1991). The stable model semantics induces not a single
 473 probability measure over interpretations, but a set of probability measures, while the
 474 well-founded semantics induces a single probability measure over assignments of
 475 three-valued logic (Cozman and Mauá 2017c). Another possibility, adopted by the P-
 476 log language, is to use the stable model semantics, but to assume a uniform probability
 477 distribution over the set of stable models of any sampled logic program (Baral et al.
 478 2009). The LP^{MLN} language also distributes probability over sets of interpretations

479 (Lee and Wang 2015), resorting to a syntax that resembles Markov logic (Sect. 6). Yet
 480 another approach is represented by a probabilistic version of the Datalog[±] language
 481 (Ceylan et al. 2016), where the stable semantics is adopted but program repairs are
 482 automatically invoked when the program is inconsistent.

483 There is no consensus yet on how to handle probabilistic logic programs with
 484 disjunctive heads and other common constructs (Cozman and Mauá 2017c). Logic
 485 Programs with Annotated Disjunctions (LPADs) offer syntax to handle disjunction
 486 probabilistically (Vennekens et al. 2004), as also done in CP-logic (Vennekens et al.
 487 2009). A LPAD may contain a rule such as

```
488     (heads(Coin):0.6) ; (tails(Coin):0.4) :- toss(Coin).
```

489 to model a biased coin. An alternative is to adapt the stable model semantics of
 490 disjunctive programs to handle probabilities (Cozman and Mauá 2017a).

491 Even more general combinations of answer set programming, first-order sen-
 492 tences, and probabilities have been investigated (Nickles and Mileo 2015); another
 493 extension that has received attention is the specification of continuous random vari-
 494 ables (Nitti et al. 2016).

495 All of these extensions of Sato’s semantics still aim at defining a distribution
 496 over interpretations of grounded atoms. A few years later than Sato’s distribution
 497 semantics, a different approach to probabilistic logic programming was advocated,
 498 with the combination of Inductive Logic Programming and probabilities (Muggleton
 499 1996). The resulting mixture has been pursued in many guises, usually under the
 500 banner of *Probabilistic Inductive Logic Programming* (De Raedt et al. 2010; De
 501 Raedt 2008). Assumptions are often distinct from Sato’s, for instance by placing
 502 probabilities over the set of proofs of a logic program (sometimes referred to as
 503 *proof theoretic semantics*) (Cussens 1999; De Raedt and Kersting 2004).

504 Two additional formalisms that specify Bayesian networks using logic program-
 505 ming principles, but not all syntactic conventions of logic programming, are Logical
 506 Bayesian Networks (LBNs) (Fierens et al. 2005, 2004) and Bayesian Logic Programs
 507 (BLPs) (De Raedt and Kersting 2004; Kersting et al. 2000). Both distinguish between
 508 guard parvariables and parvariables associated with probabilities. To illustrate, con-
 509 sider the following Bayesian Logic Program fragment (De Raedt and Kersting 2004):

```
510     carrier(x) | founder(x)
511     carrier(x) | mother(m,x), carrier(m),
512                   father(f,x), carrier(f)
```

513 where *mother* and *father* are *logical predicates* (that is, guard parvariables), while
 514 *founder* and *carrier* are associated with probabilities that must be given separately.
 515 Logical Bayesian Networks adopt additional elements of logic programming in the
 516 manipulation of logical predicates (Fierens et al. 2005).

5 Probabilistic Logic, Again; and Probabilistic Description Logics

A natural idea would be the mix graphs as used in Bayesian networks with general assessments as allowed in probabilistic logics. This is exactly the scheme adopted in Andersen and Hooker’s Bayesian Logic and its variants (Andersen and Hooker 1994; de Campos et al. 2009; Cozman et al. 2008). The resulting probabilistic logics can often be viewed as specification languages for *credal networks*; that is, for graph-based models that encode *sets* of probability distributions rather than a single distribution (Cozman 2000).

Another natural idea is to specify Bayesian networks using extended first-order formulas (Bacchus 1993). For instance, one might write

$$\forall \chi, y : \text{fails}(\chi, y) \equiv (\text{isDedicated}(\chi) \wedge \mathbf{a1}(\chi, y)) \vee (\text{isHard}(y) \wedge \mathbf{a2}(\chi, y)),$$

$$\forall \chi, y : \mathbb{P}(\mathbf{a1}(\chi, y) = 1) = 0.2, \quad \forall \chi, y : \mathbb{P}(\mathbf{a2}(\chi, y) = 1) = 0.6, \quad (1)$$

where the symbol \equiv is used to emphasize that we have a definition, and both $\mathbf{a1}$ and $\mathbf{a2}$ are auxiliary predicates (similar to the auxiliary probabilistic facts used in Example 4). It is worth noting that, for fixed χ and y , Expression (1) specifies a Noisy-OR gate for *fails* (Pearl 1988).

Expression (1) suggests a general strategy, where a Bayesian network is specified by associating each random variable X_i with a definition $X_i \equiv f(\text{pa}(X_i), Y_1, \dots, Y_m)$, for a deterministic function f of the parents of X_i and some auxiliary random variables Y_1, \dots, Y_m that are assigned probabilities $\mathbb{P}(Y_j = 1) = \alpha_j$. Poole refers to the latter random variables as *independent choices*, as he argues in favor of this specification strategy (Poole 2010). This sort of strategy is used in structural models to handle continuous random variables (Pearl 2009).

Another popular strategy in combining logic and probabilities is to mix Bayesian networks and *description logics*. Recall that description logics are (usually) decidable fragments of first-order logic that suffice for many knowledge representation tasks (Baader et al. 2017), as discussed in chapter “Reasoning with Ontologies” of Volume 1. A description logic (usually) contains a vocabulary containing *individuals*, *concepts*, and *roles*. New concepts can be defined by *intersection*, *union*, or *complement* of other concepts. And a concept can be defined using constructs $\exists r.C$ or $\forall r.C$, where r is a role and C is a concept. An inclusion axiom is written $C \sqsubseteq D$, and a definition is written $C \equiv D$, where C and D are concepts. A semantics for a description logic can (usually) be given by translation of individuals into constants, concepts into unary predicates, roles into binary predicates, and by translation of various operators into logical operators (Borgida 1996). That is, we have a domain \mathcal{D} that again is a set, and for each concept C we have a set $\mathbb{I}(C)$ of elements of the domain, and for each role r we have a set $\mathbb{I}(r)$ of pairs of elements of the domain. Inclusion $C \sqsubseteq D$ means that $\mathbb{I}(C) \subseteq \mathbb{I}(D)$, and $C \equiv D$ means that $\mathbb{I}(C) = \mathbb{I}(D)$. Moreover, the intersection $C \sqcap D$, for concepts C and D , is translated into $C(\chi) \wedge D(\chi)$; similarly union is translated into disjunction and complement is translated into negation.

556 A more involved translation is needed for $\exists r.C$; this concept should be interpreted
 557 as “the set of all x such that there is y satisfying $r(x, y)$ and $C(y)$ ”. For instance, the
 558 expression

559 $\text{Student} \sqcap \exists \text{hasChild.Female}$

560 defines the set of students who have at least a daughter. In symbols, $\exists r.C$ is interpreted
 561 as $\{x \in \mathcal{D} \mid \exists y \in \mathcal{D} : (x, y) \in \mathbb{I}(r) \wedge y \in \mathbb{I}(C)\}$. Similarly, $\forall r.C$ is interpreted as $\{x \in$
 562 $\mathcal{D} \mid \forall y \in \mathcal{D} : (x, y) \in \mathbb{I}(r) \rightarrow y \in \mathbb{I}(C)\}$.

563 An early partnership between Bayesian networks and description logics is the
 564 P-CLASSIC language: there the description logic CLASSIC is enlarged with prob-
 565 abilities in a domain-based semantics (Koller et al. 1997a). One must draw a graph
 566 where nodes refer either to concepts or to parvariables that indicate how many indi-
 567 viduals are related by roles. Inference in P-CLASSIC produces, for instance, “the
 568 probability that a randomly selected individual is a student and has a daughter”; in
 569 symbols: $\mathbb{P}(\text{Student} \sqcap \exists \text{hasChild.Female})$. Other combinations of Bayesian net-
 570 works and description logics have produced tools that resemble P-CLASSIC (Ding
 571 et al. 2006; Staker 2002; Yelland 1999).

572 An interpretation-based semantics is adopted by PR-OWL (Carvalho et al. 2013;
 573 Costa and Laskey 2006), a language that combines constructs from description log-
 574 ics with Multi-Entity Bayesian Networks (Sect. 3). Another popular strategy has
 575 been to use versions of Poole’s independent choices together with description logics
 576 (Lukasiewicz et al. 2011; Riguzzi et al. 2015); for instance, in DISPONTE one can
 577 write

578 $0.9::\text{Bird} \sqsubseteq \text{Flies}$,

579 thus mixing Problog’s syntax with an inclusion axiom (Riguzzi et al. 2015). Some
 580 languages even employ Bayesian networks to specify probabilistic choices (Ceylan
 581 and Peñaloza 2014; d’Amato et al. 2008). For instance, in the Bayesian DL-Lite \mathcal{R}
 582 language one can write

583 $\text{Bird} \sqsubseteq \text{Flies} : \text{condition} = \text{true}$

584 to mean that the inclusion axiom is present with probability $\mathbb{P}(\text{condition} = \text{true})$
 585 that is in turn given by a Bayesian network (d’Amato et al. 2008).

586 There are many other probabilistic description logics that go beyond Bayesian
 587 networks, often including sophisticated logical operators and interval-valued assess-
 588 ments (Klinov and Parsia 2011; Lukasiewicz 2008). These languages have found
 589 application in the semantic web, in information retrieval, and in knowledge rep-
 590 resentation by ontologies. The reader can benefit from substantial surveys in the
 591 literature (Lukasiewicz and Straccia 2008; Predoiu and Stuckenschmidt 2009).

6 Markov Random Fields: Undirected Graphs

We have so far considered languages where probabilities are specified directly; that is, if the probability that John is dedicated is 0.9, one just writes $\mathbb{P}(\text{JohnIsDedicated} = 1) = 0.9$. An entirely different strategy is employed in Markov random fields (Koller and Friedman 2009). There one uses an undirected graph, where each node is a random variable, and where each clique C_i of the graph is associated with a positive function $f_i(C_i)$ (recall that a clique is a complete sub-graph); see also chapter “Belief Graphical Models for Uncertainty Representation and Reasoning” of this Volume. For instance, the Markov random field defined in Fig. 8 has nine cliques containing two random variables, and two cliques containing three random variables. With rather general assumptions, the joint distribution then factorizes as

$$\mathbb{P}(X_1 = x_1, \dots, X_n = x_n) = (1/\nu) \prod_i f_i(C_i = c_i),$$

where the product goes over the set of cliques, and for each clique C_i the symbol c_i is the projection of $\{X_1 = x_1, \dots, X_n = x_n\}$ on the random variables in C_i . The normalization constant ν is called the *partition function*, and it is equal to $\sum_{x_1, \dots, x_n} \prod_i f_i(C_i = c_i)$ (see also chapter “Belief Graphical Models for Uncertainty Representation and Reasoning” of this Volume). There are technical complications if we want to find a Markov condition for such graphs in the presence of zero probabilities; we do not dwell on such details here.

The graph in Fig. 8 induces a very structured Markov random field, as we can expect the function for clique `relaxed(Ann) – smokes(Ann)` to be equal to the function for clique `relaxed(Bob) – smokes(Bob)`, and so on. It is thus natural to think about languages that would allow us to specify repetitive Markov random fields. We mention two approaches: relational Markov networks and Markov logic networks.

A *relational Markov network* is the exact counterpart of PRMs in the context of undirected graphs (Taskar et al. 2007). A relational Markov network consists, first, of a set of clique templates; each template contains parvariables and conditions on parvariables that determine which instantiations are connected by




Fig. 8 The Markov random field about smoking

edges (some implementations specify such conditions using database queries written in SQL). Given a domain (or a database), a template is instantiated into a set of random variables and edges between those random variables. The second part of a relational Markov network is a set of functions that map parvariables to positive numbers; we refer to these functions again as *parfactors* (note that these parfactors are not restricted to values in $[0, 1]$). As an example, we could build the Markov random field in Fig. 8 by specifying parvariables *relaxed*, *smokes*, and *influences*, and parfactors $f(\text{relaxed}(\chi), \text{smokes}(\chi))$ and $g(\text{smokes}(\chi), \text{influences}(\chi, y), \text{smokes}(y))$. For instance, we might specify the first parfactor as follows:

$\text{relaxed}(\chi)$	$\text{smokes}(\chi)$	$f(\text{relaxed}(\chi), \text{smokes}(\chi))$
0	0	2
0	1	1
1	0	1
1	1	3

Given an instantiation of the parvariables, the parfactors can be instantiated as well, and the result is a Markov random field such as the one in Fig. 8.

Because relational Markov networks can represent uncertainty about links, they have been successful in tasks such as collective information extraction (Bunescu and Mooney 2004) and activity recognition (Liao et al. 2006).

Markov logic networks adopt a different, albeit related, strategy (Domingos and Lowd 2009; Richardson and Domingos 2006). Instead of parvariables and parfactors, we really start with predicates and first-order sentences; the only departure from dry first-order logic is that we associate a *weight* with each sentence.

In our “smoking” example, we might contemplate two weighted sentences:

$$\begin{aligned}
 \text{Weight} = 2 : & \quad \forall \chi : \text{relaxed}(\chi) \rightarrow \text{smokes}(\chi), \\
 \text{Weight} = 4 : & \quad \forall \chi, y : \text{influences}(\chi, y) \wedge \text{smokes}(\chi) \rightarrow \text{smokes}(y).
 \end{aligned} \tag{2}$$

As before, predicates have a dual role as parvariables, and a grounding of a predicate works as a random variable.

The semantics of Markov logic networks can be explained operationally, by a procedure that takes a Markov logic network and a set \mathcal{D} , the domain, and produces a Markov random field. The procedure is rather simple: first, ground all sentences with respect to all groundings, repeating weights as appropriate. For instance, in the “smoking” example with domain $\{\text{Ann}, \text{Bob}\}$, we obtain

- Weight = 2 : relaxed(Ann) → smokes(Ann),
- Weight = 2 : relaxed(Bob) → smokes(Bob),
- 650 Weight = 4 : influences(Ann, Ann) ∧ smokes(Ann) → smokes(Ann),
- Weight = 4 : influences(Ann, Bob) ∧ smokes(Ann) → smokes(Bob),
- Weight = 4 : influences(Bob, Ann) ∧ smokes(Bob) → smokes(Ann),
- Weight = 4 : influences(Bob, Bob) ∧ smokes(Bob) → smokes(Bob).

651 Now read each one of these weighted propositional sentences as a function:

652 “Weight = $\alpha : \phi$ ” yields $\begin{cases} e^\alpha & \text{when } \phi \text{ is true,} \\ 1 & \text{otherwise.} \end{cases}$

653 This somewhat mysterious convention yields a single joint probability distribution
 654 for any Markov logic network with a given domain. If we denote by X_1, \dots, X_n the
 655 random variables that correspond to all groundings of all predicates, the convention
 656 is that:


657
$$\mathbb{P}(X_1 = x_1, \dots, X_n = x_n) = (1/Z) \exp \left(\sum_j w_j n_j \right),$$

658 where j runs over the weighted sentences, w_j is the weight of the j th weighted
 659 sentence, and n_j is the number of groundings of the j th sentence that are true
 660 for the configuration $\{X_1 = x_1, \dots, X_n = x_n\}$ (and Z is the normalization constant). In
 661 our “smoking” example, the Markov random field induced by Expression (2) has the
 662 undirected graph depicted in Fig. 8.

663 *Example 5* A short exercise is instructive. Consider the Markov logic network

- 664 Weight = $\ln(2) : A \wedge B,$ Weight = $\ln(10) : \neg A \wedge \neg B,$
- 665
- 666 Weight = $\ln(10) : B \wedge C(x).$

667 For a domain $\{a_1, \dots, a_N\}$, we get random variables $A, B, C(a_1), \dots, C(a_N)$. For
 668 instance, for $N = 2$ we get the following Markov random field:



669

670 The partition function is, using a bit of combinatorics,

671
$$Z = \left(10 \sum_{i=0}^N \binom{N}{i} \right) + \left(\sum_{i=0}^N 10^i \binom{N}{i} \right) + \left(\sum_{i=0}^N \binom{N}{i} \right) + \left(2 \sum_{i=0}^N 10^i \binom{N}{i} \right)$$

672
$$= 11 \times 2^N + 3 \times 11^N,$$

673 and then $\mathbb{P}(A = 1) = (2^N + 2 \times 11^N)/Z$ and $\mathbb{P}(B = 1) = 3 \times 11^N/Z$. \square

674 Markov logic networks offer an attractive combination of logic and probability; it
 675 is really a clever idea that has been vigorously applied in practical problems (Domingos
 676 and Lowd 2009; De Raedt et al. 2016). A clear advantage of such “undirected”
 677 languages is that they have no problem with cycles, as for instance PRMs do. One
 678 can only wonder if we have finally nailed down the laws of thought.

679 However, Markov random fields are much less transparent than PRMs when it
 680 comes to the meaning of the functions attached to cliques. These functions are not
 681 probabilities; rather, their normalized product yields the joint distribution. Interpreting
 682 the numbers is already a challenge in the propositional setting, but the relational
 683 setting introduces significant complications because the relative effect of functions
 684 depends on the number of groundings — this in turn depends on the size of the
 685 domain and on the arity of predicates. Consider Example 5: for $N = 1$ we have
 686 $\mathbb{P}(A = 1) = 24/55 < \mathbb{P}(A = 0)$, but as N grows, $\mathbb{P}(A = 1)$ grows as well, and we
 687 have $\lim_{N \rightarrow \infty} \mathbb{P}(A = 1) = 2/3$. In general, it is very hard to know, at design time,
 688 what the weights mean when instantiated (Poole et al. 2012; Jain et al. 2007).

689 An additional difficulty with Markov logic networks is that a material implication
 690 $A \rightarrow B$ is not really related to a conditional probability of B given A : $A \rightarrow B$ may be
 691 true with very high probability just because A is almost always false. Thus attaching
 692 a weight to a material implication says little about conditional probabilities.

693 And to close these remarks on Markov logic networks, consider the following
 694 short and startling example.²

695 *Example 6* Suppose we have a Markov logic network consisting of three weighted
 696 (propositional) sentences:

697 Weight = 2 : A , Weight = 1 : $A \rightarrow B$, Weight = 0 : $C \rightarrow B$.

698 Then the probability that A is true is 0.83481, but the probability that the sentence
 699 $C \rightarrow B$ is true is 0.864645. That is, the latter sentence has higher probability than
 700 the former, thus reversing the order of their weights! \square

701 There are other graph-based probabilistic models besides Bayesian networks and
 702 Markov random fields (Sadeghi and Lauritzen 2014), but it does not seem that they
 703 have been lifted to relational settings, with the notable exception of *Relational Depend-*
 704 *ency Networks* (Neville and Jensen 2007). A dependency network consists of a
 705 graph, possibly with cycles, where each node is a random variable and edges are
 706 bidirectional, and each node is associated with a conditional probability table speci-
 707 fying its probabilities given its neighbors in the graph (Heckerman et al. 2000). The
 708 joint distribution is not obtained in closed-form, but rather by a limiting procedure;
 709 it is thus a bit hard to relate local parameters and global behavior. Relational Depend-
 710 ency Networks employ the same toolset of PRMs to extend dependency networks:
 711 now the nodes of a graph are parvariables organized in classes, and a dependency

²This example is due to my colleague Marcelo Finger (personal communication).

712 network is obtained by grounding the parvariables and associated parfactors (Neville
713 and Jensen 2007). The ability to encode cyclic patterns has been useful in applica-
714 tions, for instance in natural language processing (Toutanova et al. 2003).

715 7 Probabilistic Programming

716 The emergence of PRMs and related formalisms during the nineties prompted some
717 researchers to explore programming languages with probabilistic constructs. Some
718 early formalisms, like RBNs, IBAL, and BLOG, illustrate important design choices.


719 Jaeger's *Relational Bayesian Networks (RBNs)* associate probabilities with pred-
720 icates so that, given a finite domain, one obtains a Bayesian network that specifies a
721 probability distribution over interpretations of the predicates. The design philosophy
722 behind this language is to have a few powerful constructs that are quite expressive
723 but that still allow for theoretical analysis. We have a vocabulary of predicates; a
724 predicate name stands also for a parvariable with the same arity. Each predicate is
725 associated with a *probability formula* (in essence, a parfactor specification). Here is
726 an example:

```
727 burglary(x) = 0.005;
728 alarm(x) = 0.95 burglary(x) + 0.01 (1-burglary(x));
729 cityAlarm = NoisyOr{0.8 alarm(x) | x; x = x};
```

730 The probability formula for **burglary** is the simplest one: a number. That probability
731 formula can be understood as: $\forall \chi : \mathbb{P}(\text{burglary}(\chi) = 1) = 0.005$. The probability
732 formula for **alarm** basically consists of arithmetic operations with parvariables; only
733 a few operations are allowed in RBNs. One could read this probability formula as
734 follows: for each χ , if **burglary**(χ) is true, then return 0.95; otherwise, return 0.01.
735 The probability formula for **cityAlarm** is more involved: it defines a *combination*
736 *expression* that is a Noisy-OR gate of all instances of **alarm**, each one of them
737 associated with a probability 0.8.

738 The syntax of combination expressions is rather intimidating. A combination
739 expression is written **comb**($F_1, \dots, F_k | y_1, \dots, y_m : \phi$), where each F_i is a probabil-
740 ity formula, each y_j is a logical variable that appears in the formulas F_i , ϕ is a formula
741 containing only logical variables, Boolean operators and equality; finally, **comb** is
742 a function that takes a multiset of numbers (a set with possibly repeated numbers)
743 and returns a number in $[0, 1]$. The semantics of combination expressions can be
744 explained operationally as follows. First, collect all logical variables that appear in
745 any F_1, \dots, F_k , and not in $\{y_1, \dots, y_m\}$. These are the *free* logical variables of the
746 combination expression. Basically, a combination expression is a function that yields
747 a number in $[0, 1]$ for each instance of the free logical variables. To compute this
748 number for a fixed instance of the free logical variables, go over all instances of the
749 logical variables y_1, \dots, y_m that satisfy ϕ . For each one of these instances, compute
750 the value of F_1, \dots, F_k . This produces a multiset of numbers in $[0, 1]$. Finally, apply

Fig. 9 The Bayesian network for the “alarm” RBN with domain {John, Mary}



751 comb to this multiset to obtain the value of the combination function for the fixed
 752 instance of the free variables. Clearly this is not a simple scheme, and it reveals the
 753 difficulty of producing a specification language with a small set of constructs that
 754 still can capture a large number of practical scenarios.

755 In any case, any RBN defines, for a given finite domain, a Bayesian network. For
 756 instance, we obtain the Bayesian network in Fig. 9 when we take the RBN specified
 757 previously and domain {John, Mary}. The conditional probability tables for this
 758 Bayesian network can be read from the RBN; the only non-trivial specification is the
 759 Noisy-OR gate for cityAlarm.

760 In the IBAL language (Pfeffer 2001), elements of functional programming are
 761 mixed with probabilistic assessments. A basic construct is a *stochastic choice*, syntactically
 762 expressed as $\text{dist}[p_1 : e_1, \dots, p_n : e_n]$, where each p_i is a number and each
 763 e_i is an expression: with probability p_i , the result of the statement is the evaluation of
 764 e_i . Useful syntactic sugar is provided by the `flip(p)` construct, yielding 1 with proba-
 765 bility p and 0 with probability $1 - p$ (the `flip` construct appears already in stochastic
 766 programs that inspired IBAL (Koller et al. 1997b)).

767 For example, this IBAL statement

```
768 alarm = (quake & flip(0.1)) | (burglary & flip(0.7));
```

769 assigns a Noisy-OR gate to `alarm`, by taking a disjunction (`|`) of conjunctions (`&`).
 770 And the following IBAL code is related to the University World:

```

771 student() = { isDedicated = flip 0.6; };
772 course() = { isHard = flip 0.4; };
773 registration(s,c) = {
774     fails = if s.isDedicated
775         then (if c.isHard then flip 0.4 else flip 0.1)
776         else (if c.isHard then flip 0.8 else flip 0.5);};
  
```

777 Note how PRM classes are encapsulated in functions, and parvariables appear as
 778 “local variables”. In IBAL one can also specify a domain and observations, and
 779 ask about probability values. Moreover, IBAL is not just a template language, as
 780 it offers many programming constructs that allow for arbitrary (Turing-complete)
 781 computations (Pfeffer 2001).

782 Another “generative” language is Milch et al.’s Bayesian Logic, referred to as
 783 BLOG (Milch et al. 2005a). A distinctive feature of BLOG is that one can place a
 784 distribution over the size of the domain; consider for instance (Wu et al. 2016):

```

785   Type House;
786   #House ~ Poisson(10);
787   random Boolean Burglary(House h) ~ BooleanDistrib(0.1);
788   random Boolean Earthquake ~ BooleanDistrib(0.002);
789   random Boolean Alarm(House h) ~
790     case [Burglary(h), Earthquake] in {
791       [false, false] -> BooleanDistrib(0.01),
792       [false, true]  -> BooleanDistrib(0.40),
793       [true, false]  -> BooleanDistrib(0.80),
794       [true, true]   -> BooleanDistrib(0.90)
795     };

```

796 The second line associates a Poisson distribution with the number of houses. A
797 language where domain size is not necessarily fixed is sometimes called an *open-*
798 *universe* language (Russell 2015). Such a language must deal with a number of
799 challenges; for instance, the need to consider infinitely many parents for a random
800 variable (Milch et al. 2005a). The flexibility of BLOG has met the needs of varied
801 practical applications (Russell 2015).

802 It is only natural to think that even more powerful specification languages can
803 be built by adding probabilistic constructs to existing programming languages. An
804 early language that adopted this strategy is CES, where probabilistic constructs are
805 added to C (Thrun 2000); that effort later led to the PTP language, whose syntax
806 augments CAML (Park et al. 2008). A similar strategy appeared in the community
807 interested in planning: existing languages, sometimes based on logic programming,
808 have been coupled with probabilities — two important examples are Probabilistic
809 PDDL (Yones and Littman 2004) and RDDDL (Sanner 2011). The latter languages
810 have been used extensively in demonstrations and in competitions, and have been
811 influential in using actions with deterministic and with uncertain effects to obtain
812 decision making with temporal effects.

813 A rather influential language that adds probabilistic constructs to a functional
814 programming language (in this case, Scheme) is Church (Goodman et al. 2008). Even
815 though the goal of Church was to study cognition, the language is heavily featured;
816 for instance, we can use the “flip” construct, plus conjunction and disjunction, to
817 define conditional probability tables as follows:

```

818   (define flu (flip 0.1))
819   (define cold (flip 0.2))
820   (define fever (or (and cold (flip 0.3))
821                    (and flu (flip 0.5))))

```

822 and we can even use recursion to define a geometric distribution:

```

823   (define (geometric p)
824     (if (flip p) 0 (+ 1 (geometric p))))

```

825 A descendant of Church is WebPPL; here the probabilistic constructs are added to
826 JavaScript. For instance, a conditional probability table is written as follows:

```
827   var flu = flip(0.1);
828   var cold = flip(0.2);
829   var fever = ((cold && flip(0.3)) || (flu && flip(0.5)));
```

830 Many other probabilistic programming languages have been proposed by adding
831 probabilistic constructs to programming languages from procedural to functional
832 persuasions (Gordon et al. 2014b; Kordjamshidi et al. 2015; Narayanan et al. 2016;
833 Mansinghka and Radul 2014; McCallum et al. 2009; Paige and Wood 2014; Pfeffer
834 2016; Wood et al. 2014). These languages offer at least “flip”-like commands, and
835 some offer declarative constructs that mimic plate models. For instance, in Haikaru
836 (Narayanan et al. 2016) one can specify a Latent Dirichlet Allocation model (Blei
837 et al. 2003) in a few lines of code, for instance specifying a plate as follows:

```
838   phi <~ plate _ of K: dirichlet(word_prior)
```

839 In some cases probabilistic languages have been proposed in connection with prob-
840 abilistic databases (Gordon et al. 2014a; Saad and Mansinghka 2016), a related
841 technology that we mention again in Sect. 8.

842 Another strategy in probabilistic programming is to add a powerful library to an
843 existing language. For instance, the Figaro toolkit offers a mature and complete prob-
844 abilistic modeler and reasoner on top of the programming language Scala, together
845 with solid methodological guidelines (Pfeffer 2016). Another powerful toolkit is
846 available in the Infer.NET project, a framework that meshes with several languages
847 (Minka et al. 2014). There are other toolkits that provide substantial probabilistic
848 programming features (Bessiere et al. 2013; Salvatier et al. 2016), and even recent
849 efforts to mix probabilistic programming with techniques from deep learning (Tran
850 et al. 2017).

851 A popular probabilistic programming language with a unique and powerful fea-
852 ture set is Stan (Carpenter et al. 2017), whose syntax looks deceptively similar to the
853 declarative BUGS language. In fact Stan offers imperative and modular program-
854 ming, loops and even recursion. For instance, in Stan a loop can assign values to a
855 variable repeatedly:

```
856   for (n in 1:N) {
857     t = inv_logit(a+x[n]*b);
858     y[n] ~ bernoulli(t);
859     total = total + x[n];
860   }
```

861 and one can define a recursive function such as

```
862   int random_fib(int n) {
```

```

863         if (n<2) return n + 0.02 * bernoulli(0.1);
864         else return random_fib(n-1) + random_fib(n-2);
865     }

```

866 There is now substantial activity in probabilistic programming, and the final word
867 on the subject is not yet written.³ Some languages include continuous distributions,
868 disintegrations, and even symbolic manipulation of probabilities. Most languages go
869 well beyond relational versions of Bayesian networks and Markov random fields, for
870 instance by including recursion.

871 The main goal of probabilistic programming within artificial intelligence has
872 been to “enable probabilistic modeling and machine learning to be accessible to the
873 working programmer, who has sufficient domain expertise, but perhaps not enough
874 expertise in probability theory or machine learning” (Gordon et al. 2014b). Probabilistic
875 programming is not only relevant to artificial intelligence, but also attracts
876 users interested in randomized algorithms and cryptography, and even quantum computing
877 (Barthe et al. 2015); in fact, several issues now investigated in probabilistic
878 programming stay closer to programming language design than to knowledge representation.
879

880 8 Inference and Learning: Some Brief Words

881 This survey focuses on the syntax and semantics of various languages that aim at
882 probabilistic modeling; the goal is to serve the reader a taste of what these languages
883 can do. Now, once a model is specified, it may be necessary to compute the probability
884 of various events; such a computation is called an *inference*. While the first
885 relational extensions of Bayesian networks worried mostly about inference (Bacchus
886 1993; Gilks et al. 1993; Poole 1993a; Wellman et al. 1992), a turning point was the
887 development of machine learning methods for PRMs around 1999 (Friedman et al.
888 1999). Most languages have been, since then, accompanied by appropriate learning
889 methods (De Raedt 2008; Getoor and Taskar 2007; Milch and Russell 2007; De
890 Raedt et al. 2016).

891 A discussion of inference and learning algorithms for all languages described
892 previously would certainly require another (even longer) survey. To keep matters at
893 a manageable size, only a few central results are mentioned in the remainder of this
894 section, mostly on inference algorithms.

895 Consider then the challenge of computing probabilities for a PRM, or a probabilistic
896 logic program, or some probabilistic program. One strategy is to take the
897 relational specification and translate it to a Bayesian network (or Markov random
898 field), and then run inference in the latter (Koller et al. 1997a; Wellman et al. 1992).
899 It may be that a particular inference does not require the whole Bayesian network;

³Lists of languages can be found at <http://probabilistic-programming.org/wiki/Home> and https://en.wikipedia.org/wiki/Probabilistic_programming_language.

900 only a sub-network may be *requisite* for the inference of interest. And even if the
 901 grounded Bayesian network is infinite, it may be possible to approximate inference
 902 with a suitable grounded sub-network (Laskey 2008; Pfeffer and Koller 2000).

903 In some cases, one may compute a desired probability without even generating this
 904 requisite grounded sub-network. Consider an example: suppose we have parvariables
 905 $X(\chi)$ and Y , associated with assessments:

$$906 \quad \forall \chi : \mathbb{P}(X(\chi) = 1) = 0.4, \quad Y = \begin{cases} 1 & \text{if } \exists \chi : \{X(\chi) = 1\}, \\ 0 & \text{otherwise.} \end{cases}$$

907 If we have a domain with N elements, then $\mathbb{P}(Y = 1) = 1 - (0.6)^N$; there is no
 908 need to generate the $N + 1$ requisite random variables. Example 5 shows that a more
 909 sophisticated combinatorial argument may be used to compute an inference without
 910 generating the requisite Markov random field.

911 Informally, *lifted inference* refers to a computation of probabilities that does not
 912 generate the requisite sub-network. Research on lifted inference started with a semi-
 913 nal paper by Poole (2003), who proposed a few combinatorial operators to handle
 914 some important cases. The number of lifted operators and algorithms has grown
 915 enormously, and existing surveys convey extensive references (Kersting 2012; De
 916 Raedt et al. 2016; Van den Broeck and Suciú 2017). Research on lifted inference has
 917 emphasized an abstract framework where the “input language” is just a set of par-
 918 factors (de Salvo Braz et al. 2007), similar to the ones adopted in relational Markov
 919 networks. Several techniques of lifted inference have also been applied to proba-
 920 bilistic logic programming (Riguzzi et al. 2017). One can also consider algorithms
 921 for approximate inference that look at parvariables rather than grounded random
 922 variables, thus working in a lifted fashion.

923 It is likely that most applied work must resort to approximate inference, both
 924 based on variational methods or on sampling algorithms (Koller and Friedman 2009).
 925 In fact, plate models were originally built to specify models in BUGS, a package
 926 focused on Gibbs sampling (Gilks et al. 1993). For many languages, approximate
 927 inference is all that one can hope, and the survey papers already mentioned contain a
 928 substantial number of relevant references. The use of sampling methods is particularly
 929 important in probabilistic programming, where exact inference is very difficult.
 930 Most probabilistic programming languages in essence do inference by repeatedly
 931 running code and storing statistics.

932 During the development of lifted inference, an important connection has surfaced
 933 between probabilistic logic programming and probabilistic databases. The latter consist
 934 of databases where data can be annotated with probabilities; querying such a
 935 database raises the same computational problems as running inference with parfac-
 936 tors (Suciú et al. 2011). The literature on probabilistic databases has produced deep
 937 results on the classes of queries that can be actually be solved in polynomial time.

938 A natural question then is whether one can, for a fixed a language, answer any
 939 inference in polynomial time. If a language guarantees polynomial time inference,
 940 when parvariables, parfactors and query are fixed, and the input consists of the

941 domain, then the language is *domain-liftable* (Van den Broeck 2011). While pio-
942 neering results by Jaeger show that even some rather simple languages fail to be
943 domain-liftable (Jaeger 2000, 2014), the work on lifted inference has demonstrated
944 that several important languages are domain-liftable (Taghipour et al. 2013; Van den
945 Broeck et al. 2014).

946 Domain-liftability has a narrow focus: it concentrates on polynomial complexity
947 when the input is just the domain. One can study the complexity of inference in a
948 broader framework. First, we may have, as input, the parvariables and parfactors
949 and domains and query: we then have the *inferential complexity* of the language.
950 Another possibility is to have, as input, the domains and query (the parvariables and
951 parfactors are fixed): we then have the *query complexity* of the language. Finally
952 we may have, as input, just the query (the parvariables, parfactors and domains
953 are fixed): we then have the *domain complexity* of the language. There is now a
954 substantial set of results on these notions of complexity, both for languages based
955 on first-order logic (Cozman and Mauá 2015; Mauá and Cozman 2016) and for
956 probabilistic logic programming (Cozman and Mauá 2017b, c). Besides complexity
957 questions, the theory of probabilistic programming has produced serious analysis of
958 programming patterns and semantic foundations (Gordon et al. 2014b).

959 Lifted inference is not the only genuinely relational aspect of inference for the
960 languages we have surveyed. Another important problem is deciding whether a prob-
961 abilistic program has global semantics; that is, deciding whether all possible ground-
962 ings of a probabilistic program do define a probability distribution (Jaeger 2002).
963 Yet another problem is *referential uncertainty*, where one is uncertain about an asso-
964 ciation between individuals in a domain (Getoor et al. 2007). Inference in relational
965 probabilistic languages covers a large range of techniques.

966 An even more bewildering variety of techniques has been developed in connection
967 with machine learning for languages discussed in this survey. In fact, most of those
968 languages have been proposed already with corresponding learning algorithms. This
969 is certainly true for Probabilistic Inductive Logic Programming (De Raedt 2008; De
970 Raedt et al. 2010), and indeed for most work on probabilistic logic programming
971 (Riguzzi et al. 2014; Sato 1995). Machine learning has also been a central concern
972 of probabilistic programming (Goodman et al. 2008; Gordon et al. 2014b), and a
973 basic feature of PRMs from their inception (Friedman et al. 1999; Getoor and Taskar
974 2007).

975 Broadly speaking, there are two distinct problems that learning algorithms try to
976 solve: one is *parameter learning*, where a set of syntactically correct statements is
977 given, and the exact probability values must be estimated from data; the other problem
978 is *structure learning*, where both the statements and the probability values must
979 be extracted from data. Methods based on Inductive Logic Programming typically
980 assume that a set of positive and negative examples is available, and the goal is to
981 guarantee that positive examples are derived by the learned model, while negative
982 examples are not (De Raedt 2008). Statistical methods have different assumptions,
983 usually focusing on maximization of some score that depends on the data and the
984 model. For parameter learning, the most popular score is the likelihood function;
985 for structure learning, it is necessary to penalize the complexity of the statements,

986 and a multitude of scores is available, most of them imported from the literature
 987 on Bayesian networks (Koller and Friedman 2009). The main difficulty in structure
 988 learning is the size of the space of possible sets of statements; generally some heuristic
 989 search is employed (Getoor and Taskar 2007; De Raedt et al. 2016). The vast literature
 990 on data mining and machine learning offers many possible techniques to apply, as
 991 the reader can appreciate by reading the papers cited in this survey.

992 9 Conclusion

993 This chapter has visited many different languages that aim at probabilistic modeling
 994 in knowledge representation and machine learning. As noted previously, the focus
 995 of this survey has been to provide a gentle discussion of syntactic and semantic
 996 features of existing languages, without too much technical detail; a careful review
 997 of inference has not been attempted, and the broad topic of learning algorithms has
 998 been barely scratched. Moreover, we have focused on specification languages with
 999 an “artificial intelligence motivation”, and avoided languages that address specific
 1000 tasks such as risk analysis or information retrieval. A hopefully satisfactory way to
 1001 continue the study of languages for probabilistic modeling in artificial intelligence
 1002 is to consult the list of references for this chapter.

1003 It seems fair to say that there is no single language that can serve all purposes; a
 1004 skilled practitioner cannot hope to use a single language in every application. One
 1005 can find languages that are narrow and efficient, and languages that are flexible and
 1006 that resist exact inference. One can find declarative and imperative languages, with
 1007 both pros and cons of these programming paradigms. And certainly the future will
 1008 bring an even more diverse zoo of languages. No unified set of Laws of Thought
 1009 emerges from the current literature.

1010 **Acknowledgements** The author was partially supported by CNPq (grant 308433/2014-9). This
 1011 work was partially supported by FAPESP (grant 2016/18841-0).

1012 References

- 1013 Abadi M, Halpern JY (1994) Decidability and expressiveness for first-order logics of probability.
 1014 *Inf Comput* 112(1):1–36
- 1015 Andersen KA, Hooker JN (1994) Bayesian logic. *Decis Support Syst* 11:191–210
- 1016 Baader F, Nutt W (2002) Basic description logics. *Description logic handbook*. Cambridge Univer-
 1017 sity Press, Cambridge, pp 47–100
- 1018 Baader F, Horrocks I, Lutz C, Sattler U (2017) *An introduction to description logic*. Cambridge
 1019 University Press, Cambridge
- 1020 Bacchus F (1990) *Representing and reasoning with probabilistic knowledge: a logical approach*.
 1021 MIT Press, Cambridge
- 1022 Bacchus F (1993) Using first-order probability logic for the construction of Bayesian networks. In:
 1023 *Conference on uncertainty in artificial intelligence*, pp 219–226

- 1024 Baral C (2003) Knowledge representation, reasoning, and declarative problem solving. Cambridge
1025 University Press, Cambridge
- 1026 Baral C, Gelfond M, Rushton N (2009) Probabilistic reasoning with answer sets. *Theory Pract Log*
1027 *Program* 9(1):57–144
- 1028 Barthe G, Gordon AD, Katoen JP, McIver A (2015) Challenges and trends in probabilistic program-
1029 ming. In: Dagstuhl reports (Seminar 15181), vol 5. Dagstuhl Publishing, pp 123–141
- 1030 Bessiere P, Mazer E, Ahuactzin JM, Mekhnacha K (2013) Bayesian programming. CRC Press,
1031 Boca Raton
- 1032 Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. *J Mach Learn Res* 3:993–1022
- 1033 Boole G (1958) The laws of thought. Dover edition, New York
- 1034 Borgida A (1996) On the relative expressiveness of description logics and predicate logics. *Artif*
1035 *Intell* 82(1–2):353–367
- 1036 Bruno G, Gilio A (1980) Applicazione del metodo del simplesso al teorema fondamentale per le
1037 probabilità nella concezione soggettivistica. *Statistica* 40:337–344
- 1038 Bunescu R, Mooney RJ (2004) Collective information extraction with relational Markov networks.
1039 In: Annual meeting of the association for computational linguistics
- 1040 Buntine WL (1994) Operations for learning with graphical models. *J Artif Intell Res* 2:159–225
- 1041 Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li
1042 P, Riddell A (2017) Stan: a probabilistic programming language. *J Stat Softw* 76(1):1–32. <https://doi.org/10.18637/jss.v076.i01>
- 1044 Carvalho RN, Laskey KB, Costa PCG, Ladeira M, Santos LL, Matsumoto S (2010) UnBayes:
1045 modeling uncertainty for plausible reasoning in the semantic web. In: Wu G (ed) *Semantic web*.
1046 InTech, pp 1–28
- 1047 Carvalho RN, Laskey KB, Costa PC (2013) PR-OWL 2.0 — bridging the gap to OWL semantics.
1048 In: URSW 2008-2010/UniDL 2010, LNAI 7123. Springer, pp 1–18
- 1049 Ceylan ÍÍ, Peñaloza R (2014) The Bayesian description logic \mathcal{BEL} . In: International joint conference
1050 on automated reasoning, pp 480–494
- 1051 Ceylan ÍÍ, Lukasiewicz T, Peñaloza R (2016) Complexity results for probabilistic Datalog $^{\pm}$. In:
1052 European conference on artificial intelligence, pp 1414–1422
- 1053 Clark KL (1978) Negation as failure. *Logic and data bases*. Springer, Berlin, pp 293–322
- 1054 Coletti G, Scozzafava R (2002) Probabilistic logic in a coherent setting. In: *Trends in logic*, vol 15.
1055 Kluwer, Dordrecht
- 1056 Costa PCG, Laskey KB (2006) PR-OWL: a framework for probabilistic ontologies. In: *Conference*
1057 *on formal ontology in information systems*
- 1058 Costa VS, Page D, Qazi M, Cussens J (2003) CLP(\mathcal{BN}): constraint logic programming for proba-
1059 bilistic knowledge. In: Kjaerulff U, Meek C (eds) *Conference on uncertainty in artificial intelli-*
1060 *gence*. Morgan-Kaufmann, San Francisco, pp 517–524
- 1061 Cozman FG (2000) Credal networks. *Artif Intell* 120:199–233
- 1062 Cozman FG, Mauá DD (2015) The complexity of plate probabilistic models. In: *Scalable uncertainty*
1063 *management*. LNCS, vol 9310. Springer, Cham, pp 36–49
- 1064 Cozman FG, Mauá DD (2017a) The complexity of inferences and explanations in probabilistic logic
1065 programming. *Symbolic and quantitative approaches to reasoning with uncertainty*. Lecture notes
1066 in computer science, vol 10369. Springer, Cham, pp 449–458
- 1067 Cozman FG, Mauá DD (2017b) On the complexity of propositional and relational credal networks.
1068 *Int J Approx Reason* 83:298–319
- 1069 Cozman FG, Mauá DD (2017c) On the semantics and complexity of probabilistic logic programs.
1070 *J Artif Intell Res* 60:221–262
- 1071 Cozman FG, de Campos CP, da Rocha JCF (2008) Probabilistic logic with independence. *Int J*
1072 *Approx Reason* 49:3–17
- 1073 Cussens J (1999) Parameter estimation in stochastic logic programs. *Mach Learn* 44(3):245–271
- 1074 da Costa PCG, Laskey KB (2005) Of Klingons and starships: Bayesian logic for the 23rd century.
1075 In: *Conference on uncertainty in artificial intelligence*

- 1076 d'Amato C, Fanizzi N, Lukasiewicz T (2008) Tractable reasoning with Bayesian description logics.
 1077 In: Greco S, Lukasiewicz T (eds) International conference on scalable uncertainty management.
 1078 Lecture notes in computer science, vol 5291. Springer, pp 146–159
- 1079 Dantsin E, Eiter T, Voronkov A (2001) Complexity and expressive power of logic programming.
 1080 ACM Comput Surv 33(3):374–425
- 1081 Darwiche A (2009) Modeling and reasoning with Bayesian networks. Cambridge University Press,
 1082 Cambridge
- 1083 De Bona G, Cozman FG (2017) Encoding the consistency of relational Bayesian networks. In:
 1084 Encontro Nacional de Inteligência Artificial e Computacional, Uberlândia, Brasil
- 1085 de Campos CP, Cozman FG, Luna JEO (2009) Assembling a consistent set of sentences in relational
 1086 probabilistic logic with stochastic independence. J Appl Log 7:137–154
- 1087 de Finetti B (1964) Foresight: its logical laws, its subjective sources. In: Kyburg HE Jr, Smokler
 1088 HE (eds) Studies in subjective probability. Wiley, New York
- 1089 De Raedt L (2008) Logical and relational learning. Springer, Berlin
- 1090 De Raedt L, Kersting K (2004) Probabilistic inductive logic programming. In: International confer-
 1091 ence on algorithmic learning theory, pp 19–36
- 1092 De Raedt LD, Kimmig A (2015) Probabilistic (logic) programming concepts. Mach Learn 100:5–47
- 1093 De Raedt L, Frasconi P, Kersting K, Muggleton S (2010) Probabilistic inductive logic programming.
 1094 Springer, Berlin
- 1095 De Raedt LD, Kersting K, Natarajan S, Poole D (2016) Statistical relational artificial intelligence:
 1096 logic, probability, and computation. Morgan & Claypool Publishers, San Rafael
- 1097 de Salvo Braz R, Amir E, Roth D (2007) Lifted first-order probabilistic inference. In: Getoor L,
 1098 Taskar B (eds) An introduction to statistical relational learning. MIT Press, Cambridge, pp 433–
 1099 451
- 1100 Ding Z, Peng Y, Pan R (2006) BayesOWL: uncertainty modeling in semantic web ontologies. In:
 1101 Soft computing in ontologies and semantic web. Studies in fuzziness and soft computing, vol
 1102 204. Springer, Berlin, pp 3–29
- 1103 Domingos P, Lowd D (2009) Markov logic: an interface layer for artificial intelligence. Morgan
 1104 and Claypool, San Rafael
- 1105 Eiter T, Ianni G, Krennwalner T (2009) Answer set programming: a primer. Reasoning web.
 1106 Springer, Berlin, pp 40–110
- 1107 Enderton HB (1972) A mathematical introduction to logic. Academic, Orlando
- 1108 Fagin R, Halpern JY, Megiddo N (1990) A logic for reasoning about probabilities. Inf Comput
 1109 87:78–128
- 1110 Fierens D, Blockeel H, Ramon J, Bruynooghe M (2004) Logical Bayesian networks. In: Workshop
 1111 on multi-relational data mining, pp 19–30
- 1112 Fierens D, Blockeel H, Bruynooghe M, Ramon J (2005) Logical Bayesian networks and their
 1113 relation to other probabilistic logical models. In: Conference on inductive logic programming,
 1114 pp 121–135
- 1115 Fierens D, Van den Broeck G, Renkens J, Shrerionov D, Gutmann B, Janssens G, De Raedt L (2014)
 1116 Inference and learning in probabilistic logic programs using weighted Boolean formulas. Theory
 1117 Pract Log Program 15(3):358–401
- 1118 Friedman N, Getoor L, Koller D, Pfeffer A (1999) Learning probabilistic relational models. In:
 1119 International joint conference on artificial intelligence, pp 1300–1309
- 1120 Fuhr N (1995) Probabilistic datalog - a logic for powerful retrieval methods. Conference on research
 1121 and development in information retrieval, Seattle, Washington, pp 282–290
- 1122 Gaifman H (1964) Concerning measures on first-order calculi. Isr J Math 2:1–18
- 1123 Gaifman H, Snir M (1982) Probabilities over rich languages, testing and randomness. J Symb Log
 1124 47(3):495–548
- 1125 Gelfond M, Lifschitz V (1988) The stable model semantics for logic programming. Proceedings of
 1126 international logic programming conference and symposium 88:1070–1080
- 1127 Getoor L, Grant J (2006) PRL: a probabilistic relational language. Mach Learn 62:7–31
- 1128 Getoor L, Taskar B (2007) Introduction to statistical relational learning. MIT Press, Cambridge

- 1129 Getoor L, Friedman N, Koller D, Pfeffer A, Taskar B (2007) Probabilistic relational models. In:
1130 Introduction to statistical relational learning, MIT Press, Cambridge
- 1131 Gilks WR, Thomas A, Spiegelhalter D (1993) A language and program for complex Bayesian
1132 modelling. *The Statistician* 43:169–178
- 1133 Glesner S, Koller D (1995) Constructing flexible dynamic belief networks from first-order proba-
1134 bilistic knowledge bases. In: Symbolic and quantitative approaches to reasoning with uncertainty,
1135 pp 217–226
- 1136 Goldman RP, Charniak E (1990) Dynamic construction of belief networks. In: Conference of uncer-
1137 tainty in artificial intelligence, pp 90–97
- 1138 Goodman ND, Mansinghka VK, Roy D, Bonawitz K, Tenenbaum JB (2008) Church: a language
1139 for generative models. In: Conference in uncertainty in artificial intelligence, pp 220–229
- 1140 Gordon AD, Grapple T, Rolland N, Russo C, Bergstrom J, Guiver J (2014a) Tabular: a schema-driven
1141 probabilistic programming language. *ACM SIGPLAN Not* 49(1):321–334
- 1142 Gordon AD, Henzinger TA, Nori AV, Rajmani SK (2014b) Probabilistic programming. In: Proceed-
1143 ings of the conference on future of software engineering. ACM, New York, pp 167–181. <https://doi.org/10.1145/2593882.2593900>
- 1144
- 1145 Hadjichristodoulou S, Warren DS (2012) Probabilistic logic programming with well-founded nega-
1146 tion. In: International symposium on multiple-valued logic, pp 232–237
- 1147 Hailperin T (1976) Boole’s logic and probability: a critical exposition from the standpoint of con-
1148 temporary algebra, logic, and probability theory. North-Holland, Amsterdam
- 1149 Hailperin T (1996) Sentential probability logic. Lehigh University Press, Bethlehem
- 1150 Halpern JY (2003) Reasoning about uncertainty. MIT Press, Cambridge
- 1151 Hansen P, Jaumard B (1996) Probabilistic satisfiability. Technical report G-96-31, Les Cahiers du
1152 GERAD, École Polytechnique de Montréal
- 1153 Heckerman D, Chickering DM, Meek C, Rounthwaite R, Kadie C (2000) Dependency networks
1154 for inference, collaborative filtering, and data visualization. *J Mach Learn Res* 1:49–75
- 1155 Heckerman D, Meek C, Koller D (2007) Probabilistic entity-relationship models, PRMs, and plate
1156 models. In: Taskar B, Getoor L (eds) Introduction to statistical relational learning. MIT Press,
1157 Cambridge, pp 201–238
- 1158 Hoover DN (1978) Probability logic. *Ann Math Log* 14:287–313
- 1159 Horsch MC, Poole D (1990) A dynamic approach to probabilistic inference using Bayesian net-
1160 works. In: Conference of uncertainty in artificial intelligence, pp 155–161
- 1161 Jaeger M (2000) On the complexity of inference about probabilistic relational models. *Artif Intell*
1162 117(2):297–308
- 1163 Jaeger M (2002) Relational Bayesian networks: a survey. *Linkop Electron Artic Comput Inf Sci* 6
- 1164 Jaeger M (2014) Lower complexity bounds for lifted inference. *Theory Pract Log Program*
1165 15(2):246–264
- 1166 Jain D, Kirchlechner B, Beetz M (2007) Extending Markov logic to model probability distributions
1167 in relational domains. In: *KI 2007: advances in artificial intelligence. Lecture Notes in Computer*
1168 *Science*, vol 4667. Springer, Berlin
- 1169 Keisler HJ (1985) Probabilistic quantifiers. In: Barwise J, Feferman S (eds) *Model-theoretic logic*.
1170 Springer, New York, pp 509–556
- 1171 Kersting K (2012) Lifted probabilistic inference. In: De Raedt L, Bessiere C, Dubois D, Doherty
1172 P, Frasconi P, Heintz F, Lucas P (eds) *European conference on artificial intelligence*. IOS Press,
1173 Amsterdam
- 1174 Kersting K, De Raedt L, Kramer S (2000) Interpreting Bayesian logic programs. In: *AAAI-2000*
1175 *workshop on learning statistical models from relational data*
- 1176 Klinov P, Parsia B (2011) A hybrid method for probabilistic satisfiability. In: Bjorner N, Sofronie-
1177 Stokkermans V (eds) *International conference on automated deduction*. Springer, Berlin, pp
1178 354–368
- 1179 Koller D, Friedman N (2009) Probabilistic graphical models: principles and techniques. MIT Press,
1180 Cambridge

- 1181 Koller D, Pfeffer A (1997) Object-oriented Bayesian networks. In: Conference on uncertainty in
 1182 artificial intelligence, pp 302–313
- 1183 Koller D, Pfeffer A (1998) Probabilistic frame-based systems. In: National conference on artificial
 1184 intelligence (AAAI), pp 580–587
- 1185 Koller D, Levy AY, Pfeffer A (1997a) P-CLASSIC: a tractable probabilistic description logic. In:
 1186 AAAI, pp 390–397
- 1187 Koller D, McAllester D, Pfeffer A (1997b) Effective Bayesian inference for stochastic programs.
 1188 In: AAAI, pp 740–747
- 1189 Kordjamshidi P, Roth D, Wu H (2015) Saul: towards declarative learning based programming. In:
 1190 International joint conference on artificial intelligence (IJCAI), pp 1844–1851
- 1191 Lakshmanan LVS, Sadri F (1994) Probabilistic deductive databases. In: Symposium on logic pro-
 1192 gramming, pp 254–268
- 1193 Laskey KB (2008) MEBN: a language for first-order Bayesian knowledge bases. *Artif Intell* 172(2–
 1194 3):140–178
- 1195 Lee J, Wang Y (2015) A probabilistic extension of the stable model semantics. In: AAAI spring
 1196 symposium on logical formalizations of commonsense reasoning, pp 96–102
- 1197 Liao L, Fox D, Kautz H (2006) Location-based activity recognition. In: Advances in neural infor-
 1198 mation processing systems, pp 787–794
- 1199 Lukasiewicz T (1998) Probabilistic logic programming. In: European conference on artificial intel-
 1200 ligence, pp 388–392
- 1201 Lukasiewicz T (2005) Probabilistic description logic programs. In: Proceedings of the 8th European
 1202 conference on symbolic and quantitative approaches to reasoning with uncertainty (ECSQARU
 1203 2005). pp 737–749, Springer, Barcelona
- 1204 Lukasiewicz T (2008) Expressive probabilistic description logics. *Artif Intell* 172(6–7):852–883
- 1205 Lukasiewicz T, Straccia U (2008) Managing uncertainty and vagueness in description logics for the
 1206 semantic web. *J Web Semant* 6:291–308
- 1207 Lukasiewicz T, Predoiu L, Stuckenschmidt H (2011) Tightly integrated probabilistic description
 1208 logic programs for representing ontology mappings. *Ann Math Artif Intell* 63(3/4):385–425
- 1209 Lunn D, Spiegelhalter D, Thomas A, Best N (2009) The BUGS project: evolution, critique and
 1210 future directions. *Stat Med* 28:3049–3067
- 1211 Lunn D, Jackson C, Best N, Thomas A, Spiegelhalter D (2012) The BUGS book: a practical
 1212 introduction to Bayesian analysis. CRC Press/Chapman and Hall, Boca Raton
- 1213 Mahoney S, Laskey KB (1996) Network engineering for complex belief networks. In: Conference
 1214 on uncertainty in artificial intelligence
- 1215 Mansinghka V, Radul A (2014) CoreVenture: a highlevel, reflective machine language for proba-
 1216 bilistic programming. In: NIPS workshop on probabilistic programming
- 1217 Mauá DD, Cozman FG (2016) The effect of combination functions on the complexity of relational
 1218 Bayesian networks. In: Conference on probabilistic graphical models — JMLR workshop and
 1219 conference proceedings, vol 52, pp 333–344
- 1220 McCallum A, Schultz K, Singh S (2009) Factorie: probabilistic programming via imperatively
 1221 defined factor graphs. In: Advances in neural information processing systems (NIPS), pp 1249–
 1222 1257
- 1223 McCarthy J, Hayes PJ (1969) Some philosophical problems from the standpoint of artificial intel-
 1224 ligence. In: Meltzer B, Michie D (eds) *Machine intelligence*, vol 4. Edinburgh University Press,
 1225 pp 463–502
- 1226 Milch B, Russell S (2007) First-order probabilistic languages: into the unknown. In: International
 1227 conference on inductive logic programming
- 1228 Milch B, Marthi B, Russell S, Sontag D, Ong DL, Kolobov A (2005a) BLOG: probabilistic models
 1229 with unknown objects. In: IJCAI
- 1230 Milch B, Marthi B, Sontag D, Russell S, Ong DL, Kolobov A (2005b) Approximate inference for
 1231 infinite contingent Bayesian networks. In: *Artificial intelligence and statistics*
- 1232 Minka T, Winn JM, Guiver JP, Webster S, Zaykov Y, Yangel B, Spengler A, Bronskill J (2014)
 1233 Infer.NET 2.6. Technical report, Microsoft Research Cambridge

- 1234 Muggleton S (1996) Stochastic logic programs. *Advances in inductive logic programming*. IOS
1235 Press, Amsterdam, pp 254–264
- 1236 Narayanan P, Carette J, Romano W, Shan C, Zinkov R (2016) Probabilistic inference by program
1237 transformation in Hakaru (system description). In: *Functional and logic programming*, pp 62–79
- 1238 Neville J, Jensen D (2007) Relational dependency networks. *J Mach Learn Res* 8:653–692
- 1239 Ng R, Subrahmanian VS (1992) Probabilistic logic programming. *Inf Comput* 101(2):150–201
- 1240 Ngo L, Haddawy P (1997) Answering queries from context-sensitive probabilistic knowledge bases.
1241 *Theor Comput Sci* 171(1–2):147–177
- 1242 Nickles M, Mileo A (2015) A system for probabilistic inductive answer set programming. In: *Inter-
1243 national Conference on scalable uncertainty management*, vol 9310. *Lecture notes in computer
1244 science*, pp 99–105
- 1245 Nilsson NJ (1986) Probabilistic logic. *Artif Intell* 28:71–87
- 1246 Nitti D, Laet TD, Raedt LD (2016) Probabilistic logic programming with hybrid domains. *Mach
1247 Learn* 103(3):407–449
- 1248 Paige B, Wood F (2014) A compilation target for probabilistic programming languages. *International
1249 conference on machine learning*, *JMLR* 32:1935–1943
- 1250 Park S, Pfenning F, Thrun S (2008) A probabilistic language based on sampling functions. *ACM
1251 Trans Program Lang Syst* 31(1):4:1–4:45
- 1252 Pearl J (1988) Probabilistic reasoning in intelligent systems: networks of plausible inference. *Morgan
1253 Kaufmann, San Mateo*
- 1254 Pearl J (2009) *Causality: models, reasoning, and inference*, 2nd edn. Cambridge University Press,
1255 Cambridge
- 1256 Pfeffer A (2001) IBAL: a probabilistic rational programming language. In: *International joint con-
1257 ference on artificial intelligence*, pp 733–740
- 1258 Pfeffer A (2016) *Practical probabilistic programming*. Manning Publications, Shelter Island
- 1259 Pfeffer A, Koller D (2000) Semantics and inference for recursive probability models. In: *AAAI*, pp
1260 538–544
- 1261 Poole D (1993a) Average-case analysis of a search algorithm for estimating prior and posterior prob-
1262 abilities in Bayesian networks with extreme probabilities. In: *13th international joint conference
1263 on artificial intelligence*, pp 606–612
- 1264 Poole D (1993b) Probabilistic Horn abduction and Bayesian networks. *Artif Intell* 64:81–129
- 1265 Poole D (1997) The independent choice logic for modelling multiple agents under uncertainty. *Artif
1266 Intell* 94(1/2):7–56
- 1267 Poole D (2003) First-order probabilistic inference. In: *International joint conference on artificial
1268 intelligence (IJCAI)*, pp 985–991
- 1269 Poole D (2008) The independent choice logic and beyond. In: De Raedt L, Frasconi P, Kersting
1270 K, Muggleton S (eds) *Probabilistic inductive logic programming*. *Lecture Notes in Computer
1271 Science*, vol 4911. Springer, Berlin, pp 222–243
- 1272 Poole D (2010) Probabilistic programming languages: independent choices and deterministic sys-
1273 tems. In: Dechter R, Geffner H, Halpern JY (eds) *Heuristics, probability and causality — a tribute
1274 to Judea pearl*. College Publications, pp 253–269
- 1275 Poole D, Buchman D, Natarajan S, Kersting K (2012) Aggregation and population growth: the
1276 relational logistic regression and Markov logic cases. In: *International Workshop on Statistical
1277 Relational AI*
- 1278 Pourret O, Naim P, Marcot B (2008) *Bayesian networks – a practical guide to applications*. Wiley,
1279 New York
- 1280 Predoiu L, Stuckenschmidt H (2009) Probabilistic models for the semantic web. The semantic web
1281 for knowledge and data management: technologies and practices. *IGI Global, Hershey*, pp 74–105
- 1282 Richardson M, Domingos P (2006) Markov logic networks. *Mach Learn* 62(1–2):107–136
- 1283 Riguzzi F (2015) The distribution semantics is well-defined for all normal programs. In: Riguzzi F,
1284 Vennkens J (eds) *International workshop on probabilistic logic programming*, *CEUR workshop
1285 proceedings*, vol 1413, pp 69–84

- 1286 Riguzzi F, Bellodi E, Zese R (2014) A history of probabilistic inductive logic programming. *Front*
 1287 *Robot AI* 1:1–5
- 1288 Riguzzi F, Bellodi E, Lamma E, Zese R (2015) Probabilistic description logics under the distribution
 1289 semantics. *Semant Web* 6(5):477–501
- 1290 Riguzzi F, Bellodi E, Zese R, Cota G, Lamma E (2017) A survey of lifted inference approaches for
 1291 probabilistic logic programming under the distribution semantics. *Int J Approx Reason* 80:313–
 1292 333
- 1293 Russell S (2015) Unifying logic and probability. *Commun ACM* 58(7):88–97
- 1294 Saad F, Mansinghka VK (2016) A probabilistic programming approach to probabilistic data analysis.
 1295 In: *Advances in neural information processing systems (NIPS)*
- 1296 Sadeghi K, Lauritzen S (2014) Markov properties for mixed graphs. *Bernoulli* 20(2):676–696
- 1297 Salvatier J, Wiecki TV, Fonnesbeck C (2016) Probabilistic programming in Python using PyMC3.
 1298 *PeerJ*
- 1299 Sanner S (2011) Relational dynamic influence diagram language (RDDL): language description.
 1300 Technical report, NICTA and Australian National University
- 1301 Sato T (1995) A statistical learning method for logic programs with distribution semantics. In:
 1302 *Conference on logic programming*, pp 715–729
- 1303 Sato T, Kameya Y (2001) Parameter learning of logic programs for symbolic-statistical modeling.
 1304 *J Artif Intell Res* 15:391–454
- 1305 Sato T, Kameya Y, Zhou NF (2005) Generative modeling with failure in PRISM. In: *International*
 1306 *joint conference on artificial intelligence*, pp 847–852
- 1307 Scott D, Krauss P (1966) Assigning probabilities to logical formulas. In: *Suppes Hintikka (ed)*
 1308 *Aspects of inductive logic*. North-Holland, Amsterdam, pp 219–264
- 1309 Staker R (2002) Reasoning in expressive description logics using belief networks. *International*
 1310 *conference on information and knowledge engineering*. Las Vegas, USA, pp 489–495
- 1311 Suciu D, Oiteanu D, Ré C, Koch C (2011) Probabilistic databases. Morgan & Claypool Publishers,
 1312 San Rafael
- 1313 Taghipour N, Fierens D, Van den Broeck G, Davis J, Blockeel H (2013) Completeness results
 1314 for lifted variable elimination. *International conference on artificial intelligence and statistics*
 1315 (AISTATS). Scottsdale, USA, pp 572–580
- 1316 Taskar B, Abbeel P, Wong MF, Koller D (2007) Relational Markov networks. In: *Getoor L, Taskar*
 1317 *B (eds) Introduction to statistical relational learning*. MIT Press, Cambridge, pp 175–199
- 1318 Thrun S (2000) Towards programming tools for robots that integrate probabilistic computation and
 1319 learning. In: *IEEE international conference on robotics and automation (ICRA)*
- 1320 Toutanova K, Klein D, Manning CD, Singer Y (2003) Feature-rich part-of-speech tagging with a
 1321 cyclic dependency network. *Conference of the North American chapter of the association for*
 1322 *computational linguistics on human language technology* 1:173–180
- 1323 Tran D, Hoffman MD, Saurous RA, Brevdo E, Murphy K, Blei DM (2017) Deep probabilistic
 1324 programming. In: *International conference on learning representations*
- 1325 Van den Broeck G (2011) On the completeness of first-order knowledge compilation for lifted
 1326 probabilistic inference. In: *Neural processing information systems*, pp 1386–1394
- 1327 Van den Broeck G, Suciu D (2017) Query processing on probabilistic data: a survey. *Found Trends*
 1328 *Databases* 7:197–341
- 1329 Van den Broeck G, Wannes M, Darwiche A (2014) Skolemization for weighted first-order model
 1330 counting. In: *International conference on principles of knowledge representation and reasoning*,
 1331 pp 111–120
- 1332 Van Gelder A, Ross KA, Schlipf JS (1991) The well-founded semantics for general logic programs.
 1333 *J Assoc Comput Mach* 38(3):620–650
- 1334 Vennekens J, Verbaeten S, Bruynooghe M (2004) Logic programs with annotated disjunctions. In:
 1335 *Logic programming - ICLP. LNCS*, vol 3132. Springer, Berlin, pp 431–445
- 1336 Vennekens J, Denecker M, Bruynooghe M (2009) CP-logic: a language of causal probabilistic
 1337 events and its relation to logic programming. *Theory Pract Log Program* 9(3):245–308

- 1338 Wellman MP, Breese JS, Goldman RP (1992) From knowledge bases to decision models. *Knowl*
1339 *Eng Rev* 7(1):35–53
- 1340 Wood F, van de Meent JW, Mansinghka V (2014) A new approach to probabilistic programming
1341 inference. In: *International conference on artificial intelligence and statistics*, pp 1024–1032
- 1342 Wu Y, Li L, Russell S, Bodik R (2016) Swift: compiled inference for probabilistic programming
1343 languages. In: *International joint conference on artificial intelligence (IJCAI)*
- 1344 Yelland PM (1999) Market analysis using a combination of Bayesian networks and description
1345 logics. Technical report SMLI TR-99-78, Sun Microsystems Laboratories
- 1346 Yones HLS, Littman ML (2004) PPDDL 1.0: an extension to PDDL for expressing planning domains
1347 with probabilistic effects. Technical report CMU-CS-04-167, Carnegie Mellon University, Pitts-
1348 burgh, PA