

Just Add Weights: Markov Logic for the Semantic Web

Pedro Domingos¹, Daniel Lowd¹, Stanley Kok¹, Hoifung Poon¹, Matthew Richardson², and Parag Singla¹

¹ Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350, U.S.A.

{pedrod, lowd, koks, hoifung, parag}@cs.washington.edu

² Microsoft Research
Redmond, WA 98052
mattri@microsoft.com

Abstract. In recent years, it has become increasingly clear that the vision of the Semantic Web requires uncertain reasoning over rich, first-order representations. Markov logic brings the power of probabilistic modeling to first-order logic by attaching weights to logical formulas and viewing them as templates for features of Markov networks. This gives natural probabilistic semantics to uncertain or even inconsistent knowledge bases with minimal engineering effort. Inference algorithms for Markov logic draw on ideas from satisfiability, Markov chain Monte Carlo and knowledge-based model construction. Learning algorithms are based on the conjugate gradient algorithm, pseudo-likelihood and inductive logic programming. Markov logic has been successfully applied to problems in entity resolution, link prediction, information extraction and others, and is the basis of the open-source *Alchemy* system.

1 Introduction

The vision of the Semantic Web is that of a web of information that computers can understand and reason about, organically built with no central organization except for a common set of standards [1]. This promises the ability to answer more complex queries and build more intelligent and effective agents than ever before. The standard languages that have been introduced so far are generally special cases of first-order logic, allowing users to define ontologies, express a rich set of relationships among objects of different types, logical dependencies between them, etc.

Fulfilling this promise, however, requires more than purely logical representations and inference algorithms. Most things in the world have some degree of uncertainty or noise – future events, such as weather and traffic, are unpredictable; information is unreliable, either from error or deceit; even simple concepts such as “fruit” and “vegetable” are imprecisely and inconsistently applied. Any system that hopes to represent varied information about the world

must therefore acknowledge the uncertain, inconsistent, and untrustworthy nature of that knowledge.

The Semantic Web project faces additional or exacerbated sources of uncertainty in a number of areas. Matching entities, ontologies and schemas is essential for linking data from different sources, but is also inherently uncertain. Moreover, data may contain false or contradictory information. To simply exclude noisy or untrusted sources is an inadequate solution since even trusted sources may have some errors and even noisy sources may have useful information to contribute. A final problem is incomplete information; when information is missing we may be able to conclude very little with certainty, but it would be a mistake to ignore the partial evidence entirely.

Markov logic is a simple yet powerful solution to the problem of integrating logic and uncertainty. Given an existing knowledge base in first-order logic, Markov logic attaches a weight to each formula. Semantically, weighted formulas are viewed as templates for constructing Markov networks. This yields a well-defined probability distribution in which worlds are more likely when they satisfy a higher-weight set of ground formulas. Intuitively, the magnitude of the weight corresponds to the relative strength of its formula; in the infinite-weight limit, Markov logic reduces to first-order logic. Since Markov logic is a direct extension of first-order logic, it does not invalidate or conflict with the existing Semantic Web infrastructure. With Markov logic, Semantic Web languages can be made probabilistic simply by adding weights to statements, and Semantic Web inference engines can be extended to perform probabilistic reasoning simply by passing the proof DAG (directed acyclic graph), with weights attached, to a probabilistic inference system. Weights may be set by hand, inferred from various sources (e.g., trust networks), or learned automatically from data. We have also developed algorithms for learning or correcting formulas from data. Markov logic has already been used to efficiently develop state-of-the-art models for entity resolution, ontology induction, information extraction, social networks, collective classification, and many other problems important to the Semantic Web. All of our algorithms, as well as sample datasets and applications, are available in the open-source Alchemy system [16] (alchemy.cs.washington.edu).

In this chapter, we describe the Markov logic representation and give an overview of current inference and learning algorithms for it. We begin with some background on first-order logic and Markov networks.

2 First-Order Logic

A *first-order knowledge base (KB)* is a set of sentences or formulas in first-order logic [10]. Formulas are constructed using four types of symbols: constants, variables, functions, and predicates. Constant symbols represent objects in the domain of interest (e.g., people: **Anna**, **Bob**, **Chris**, etc.). Variable symbols range over the objects in the domain. Function symbols (e.g., **MotherOf**) represent mappings from tuples of objects to objects. Predicate symbols represent relations among objects in the domain (e.g., **Friends**) or attributes of objects (e.g.,

Smokes). An *interpretation* specifies which objects, functions and relations in the domain are represented by which symbols. Variables and constants may be *typed*, in which case variables range only over objects of the corresponding type, and constants can only represent objects of the corresponding type. For example, the variable x might range over people (e.g., Anna, Bob, etc.), and the constant C might represent a city (e.g., Seattle, Tokyo, etc.).

A *term* is any expression representing an object in the domain. It can be a constant, a variable, or a function applied to a tuple of terms. For example, $Anna$, x , and $GreatestCommonDivisor(x, y)$ are terms. An *atomic formula* or *atom* is a predicate symbol applied to a tuple of terms (e.g., $Friends(x, MotherOf(Anna))$). Formulas are recursively constructed from atomic formulas using logical connectives and quantifiers. If F_1 and F_2 are formulas, the following are also formulas: $\neg F_1$ (negation), which is true iff F_1 is false; $F_1 \wedge F_2$ (conjunction), which is true iff both F_1 and F_2 are true; $F_1 \vee F_2$ (disjunction), which is true iff F_1 or F_2 is true; $F_1 \Rightarrow F_2$ (implication), which is true iff F_1 is false or F_2 is true; $F_1 \Leftrightarrow F_2$ (equivalence), which is true iff F_1 and F_2 have the same truth value; $\forall x F_1$ (universal quantification), which is true iff F_1 is true for every object x in the domain; and $\exists x F_1$ (existential quantification), which is true iff F_1 is true for at least one object x in the domain. Parentheses may be used to enforce precedence. A *positive literal* is an atomic formula; a *negative literal* is a negated atomic formula. The formulas in a KB are implicitly conjoined, and thus a KB can be viewed as a single large formula. A *ground term* is a term containing no variables. A *ground atom* or *ground predicate* is an atomic formula all of whose arguments are ground terms. A *possible world* (along with an interpretation) assigns a truth value to each possible ground atom.

A formula is *satisfiable* iff there exists at least one world in which it is true. The basic inference problem in first-order logic is to determine whether a knowledge base KB *entails* a formula F , i.e., if F is true in all worlds where KB is true (denoted by $KB \models F$). This is often done by *refutation*: KB entails F iff $KB \cup \neg F$ is unsatisfiable. (Thus, if a KB contains a contradiction, all formulas trivially follow from it, which makes painstaking knowledge engineering a necessity.) For automated inference, it is often convenient to convert formulas to a more regular form, typically *clausal form* (also known as *conjunctive normal form (CNF)*). A KB in clausal form is a conjunction of *clauses*, a clause being a disjunction of literals. Every KB in first-order logic can be converted to clausal form using a mechanical sequence of steps.³ Clausal form is used in resolution, a sound and refutation-complete inference procedure for first-order logic [38].

Inference in first-order logic is only semidecidable. Because of this, knowledge bases are often constructed using a restricted subset of first-order logic with more desirable properties. The two subsets most commonly applied to the Semantic Web are *Horn clauses* and *description logics*. Horn clauses are clauses containing at most one positive literal. The Prolog programming language is based on Horn

³ This conversion includes the removal of existential quantifiers by Skolemization, which is not sound in general. However, in finite domains an existentially quantified formula can simply be replaced by a disjunction of its groundings.

clause logic [21]. Prolog programs can be learned from databases by searching for Horn clauses that (approximately) hold in the data; this is studied in the field of inductive logic programming (ILP) [18]. Description logics are a decidable subset of first-order logic that is the basis of the Web Ontology Language (OWL) [7].

Table 1 shows a simple KB and its conversion to clausal form. Notice that, while these formulas may be *typically* true in the real world, they are not *always* true. In most domains it is very difficult to come up with non-trivial formulas that are always true, and such formulas capture only a fraction of the relevant knowledge. Thus, despite its expressiveness, pure first-order logic has limited applicability to practical AI problems. Many *ad hoc* extensions to address this have been proposed. In the more limited case of propositional logic, the problem is well solved by probabilistic graphical models such as Markov networks, described in the next section. We will later show how to generalize these models to the first-order case.

Table 1. Example of a first-order knowledge base and MLN. $\text{Fr}()$ is short for $\text{Friends}()$, $\text{Sm}()$ for $\text{Smokes}()$, and $\text{Ca}()$ for $\text{Cancer}()$.

First-Order Logic	Clausal Form	Weight
“Friends of friends are friends.” $\forall x \forall y \forall z \text{Fr}(x, y) \wedge \text{Fr}(y, z) \Rightarrow \text{Fr}(x, z)$	$\neg \text{Fr}(x, y) \vee \neg \text{Fr}(y, z) \vee \text{Fr}(x, z)$	0.7
“Friendless people smoke.” $\forall x (\neg(\exists y \text{Fr}(x, y)) \Rightarrow \text{Sm}(x))$	$\text{Fr}(x, g(x)) \vee \text{Sm}(x)$	2.3
“Smoking causes cancer.” $\forall x \text{Sm}(x) \Rightarrow \text{Ca}(x)$	$\neg \text{Sm}(x) \vee \text{Ca}(x)$	1.5
“If two people are friends, then either both smoke or neither does.” $\forall x \forall y \text{Fr}(x, y) \Rightarrow (\text{Sm}(x) \Leftrightarrow \text{Sm}(y))$	$\neg \text{Fr}(x, y) \vee \text{Sm}(x) \vee \neg \text{Sm}(y),$ $\neg \text{Fr}(x, y) \vee \neg \text{Sm}(x) \vee \text{Sm}(y)$	1.1 1.1

3 Markov Networks

A *Markov network* (also known as *Markov random field*) is a model for the joint distribution of a set of variables $X = (X_1, X_2, \dots, X_n) \in \mathcal{X}$ [30]. It is composed of an undirected graph G and a set of potential functions ϕ_k . The graph has a node for each variable, and the model has a potential function for each clique in the graph. A potential function is a non-negative real-valued function of the state of the corresponding clique. The joint distribution represented by a Markov network is given by

$$P(X=x) = \frac{1}{Z} \prod_k \phi_k(x_{\{k\}}) \quad (1)$$

where $x_{\{k\}}$ is the state of the k th clique (i.e., the state of the variables that appear in that clique). Z , known as the *partition function*, is given by $Z =$

$\sum_{x \in \mathcal{X}} \prod_k \phi_k(x_{\{k\}})$. Markov networks are often conveniently represented as *log-linear models*, with each clique potential replaced by an exponentiated weighted sum of features of the state, leading to

$$P(X=x) = \frac{1}{Z} \exp \left(\sum_j w_j f_j(x) \right) \quad (2)$$

A feature may be any real-valued function of the state. This chapter will focus on binary features, $f_j(x) \in \{0, 1\}$. In the most direct translation from the potential-function form (Equation 1), there is one feature corresponding to each possible state $x_{\{k\}}$ of each clique, with its weight being $\log \phi_k(x_{\{k\}})$. This representation is exponential in the size of the cliques. However, we are free to specify a much smaller number of features (e.g., logical functions of the state of the clique), allowing for a more compact representation than the potential-function form, particularly when large cliques are present. Markov logic will take advantage of this.

Inference in Markov networks is #P-complete [39]. The most widely used method for approximate inference in Markov networks is Markov chain Monte Carlo (MCMC) [11], and in particular Gibbs sampling, which proceeds by sampling each variable in turn given its Markov blanket. (The Markov blanket of a node is the minimal set of nodes that renders it independent of the remaining network; in a Markov network, this is simply the node’s neighbors in the graph.) Marginal probabilities are computed by counting over these samples; conditional probabilities are computed by running the Gibbs sampler with the conditioning variables clamped to their given values. Another popular method for inference in Markov networks is belief propagation [52].

Maximum-likelihood or MAP estimates of Markov network weights cannot be computed in closed form but, because the log-likelihood is a concave function of the weights, they can be found efficiently (modulo inference) using standard gradient-based or quasi-Newton optimization methods [28]. Another alternative is iterative scaling [8]. Features can also be learned from data, for example by greedily constructing conjunctions of atomic features [8].

4 Markov Logic

A first-order KB can be seen as a set of hard constraints on the set of possible worlds: if a world violates even one formula, it has zero probability. The basic idea in Markov logic is to soften these constraints: when a world violates one formula in the KB it is less probable, but not impossible. The fewer formulas a world violates, the more probable it is. Each formula has an associated weight (e.g., see Table 1) that reflects how strong a constraint it is: the higher the weight, the greater the difference in log probability between a world that satisfies the formula and one that does not, other things being equal.

Definition 1. [36] *A Markov logic network (MLN) L is a set of pairs (F_i, w_i) , where F_i is a formula in first-order logic and w_i is a real number. Together with*

a finite set of constants $C = \{c_1, c_2, \dots, c_{|C|}\}$, it defines a Markov network $M_{L,C}$ (Equations 1 and 2) as follows:

1. $M_{L,C}$ contains one binary node for each possible grounding of each atom appearing in L . The value of the node is 1 if the ground atom is true, and 0 otherwise.
2. $M_{L,C}$ contains one feature for each possible grounding of each formula F_i in L . The value of this feature is 1 if the ground formula is true, and 0 otherwise. The weight of the feature is the w_i associated with F_i in L .

Thus there is an edge between two nodes of $M_{L,C}$ iff the corresponding ground atoms appear together in at least one grounding of one formula in L . For example, an MLN containing the formulas $\forall x \text{Smokes}(x) \Rightarrow \text{Cancer}(x)$ (smoking causes cancer) and $\forall x \forall y \text{Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$ (friends have similar smoking habits) applied to the constants Anna and Bob (or A and B for short) yields the ground Markov network in Figure 1. Its features include $\text{Smokes}(\text{Anna}) \Rightarrow \text{Cancer}(\text{Anna})$, etc. Notice that, although the two formulas above are false as universally quantified logical statements, as weighted features of an MLN they capture valid statistical regularities, and in fact represent a standard social network model [47].




Fig. 1. Ground Markov network obtained by applying an MLN containing the formulas $\forall x \text{Smokes}(x) \Rightarrow \text{Cancer}(x)$ and $\forall x \forall y \text{Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$ to the constants Anna(A) and Bob(B).

An MLN can be viewed as a *template* for constructing Markov networks. From Definition 1 and Equations 1 and 2, the probability distribution over possible worlds x specified by the ground Markov network $M_{L,C}$ is given by

$$P(X=x) = \frac{1}{Z} \exp \left(\sum_{i=1}^F w_i n_i(x) \right) \quad (3)$$

where F is the number of formulas in the MLN and $n_i(x)$ is the number of true groundings of F_i in x . As formula weights increase, an MLN increasingly resembles a purely logical KB, becoming equivalent to one in the limit of all infinite weights. When the weights are positive and finite, and all formulas are simultaneously satisfiable, the satisfying solutions are the modes of the distribution represented by the ground Markov network.

Most importantly, Markov logic allows contradictions between formulas, which it resolves simply by weighing the evidence on both sides. This makes it well suited for merging multiple KBs. Markov logic also provides a natural and powerful approach to the problem of merging knowledge and data in different representations that do not align perfectly, as will be illustrated in the application section. Both of these tasks are also key to the success of the Semantic Web.

It is interesting to see a simple example of how Markov logic generalizes first-order logic. Consider an MLN containing the single formula $\forall \mathbf{x} \mathbf{R}(\mathbf{x}) \Rightarrow \mathbf{S}(\mathbf{x})$ with weight w , and $C = \{\mathbf{A}\}$. This leads to four possible worlds: $\{\neg \mathbf{R}(\mathbf{A}), \neg \mathbf{S}(\mathbf{A})\}$, $\{\neg \mathbf{R}(\mathbf{A}), \mathbf{S}(\mathbf{A})\}$, $\{\mathbf{R}(\mathbf{A}), \neg \mathbf{S}(\mathbf{A})\}$, and $\{\mathbf{R}(\mathbf{A}), \mathbf{S}(\mathbf{A})\}$. From Equation 3 we obtain that $P(\{\mathbf{R}(\mathbf{A}), \neg \mathbf{S}(\mathbf{A})\}) = 1/(3e^w + 1)$ and the probability of each of the other three worlds is $e^w/(3e^w + 1)$. (The denominator is the partition function Z ; see Section 3.) Thus, if $w > 0$, the effect of the MLN is to make the world that is inconsistent with $\forall \mathbf{x} \mathbf{R}(\mathbf{x}) \Rightarrow \mathbf{S}(\mathbf{x})$ less likely than the other three. From the probabilities above we obtain that $P(\mathbf{S}(\mathbf{A})|\mathbf{R}(\mathbf{A})) = 1/(1 + e^{-w})$. When $w \rightarrow \infty$, $P(\mathbf{S}(\mathbf{A})|\mathbf{R}(\mathbf{A})) \rightarrow 1$, recovering the logical entailment.

It is easily seen that all discrete probabilistic models expressible as products of potentials, including Markov networks and Bayesian networks, are expressible in Markov logic. In particular, many of the models frequently used in AI can be stated quite concisely as MLNs, and combined and extended simply by adding the corresponding formulas. Most significantly, Markov logic facilitates the construction of non-i.i.d. models (i.e., models where objects are not independent and identically distributed).

When working with Markov logic, we typically make three assumptions about the logical representation: different constants refer to different objects (unique names), the only objects in the domain are those representable using the constant and function symbols (domain closure), and the value of each function for each tuple of arguments is always a known constant (known functions). These assumptions ensure that the number of possible worlds is finite and that the Markov logic network will give a well-defined probability distribution. These assumptions are quite reasonable in most practical applications, and greatly simplify the use of MLNs. After describing how each one can be relaxed, we will make these assumptions for the remainder of the chapter. See Richardson and Domingos [36] for further details on the Markov logic representation.

The unique names assumption can be removed by introducing the equality predicate (`Equals(x, y)`, or $\mathbf{x} = \mathbf{y}$ for short) and adding the necessary axioms to the MLN: equality is reflexive, symmetric and transitive; for each unary predicate P , $\forall \mathbf{x} \forall \mathbf{y} \mathbf{x} = \mathbf{y} \Rightarrow (P(\mathbf{x}) \Leftrightarrow P(\mathbf{y}))$; and similarly for higher-order predicates and

functions [10]. This allows us to deal with instance and reference uncertainty, as illustrated in Section 7.1.

We can relax the domain closure assumption by introducing new constants to represent unknown objects. This works for any domain whose size is bounded. Markov logic can also be applied to a number of interesting infinite domains, such as when each node in the resulting infinite Markov network has a finite number of neighbors. See Singla and Domingos [43] for details on Markov logic in infinite domains.

Infinite domains can also be approximated as finite ones. Consider the transitive, anti-symmetric relation `AncestorOf(x, y)`, meaning “ x is an ancestor of y .” In a logical KB, the rule “Everyone has an ancestor” is only valid in infinite or empty domains. In Markov logic, the rule can easily be applied to finite domains, so that worlds are more likely when more objects have an ancestor within the domain. Therefore, although Markov logic semantics are well-defined for many infinite domains, a finite approach suffices for most practical applications.

Let $H_{L,C}$ be the set of all ground terms constructible from the function symbols in L and the constants in L and C (the “Herbrand universe” of (L, C)). We can remove the known function assumption by treating each element of $H_{L,C}$ as an additional constant and applying the same procedure used to remove the unique names assumption. For example, with a function $G(x)$ and constants A and B , the MLN will now contain nodes for $G(A) = A$, $G(A) = B$, etc. This leads to an infinite number of new constants, requiring the corresponding extension of MLNs. However, if we restrict the level of nesting to some maximum, the resulting MLN is still finite.

5 Inference

Recall that an MLN acts as a template for a Markov network. Therefore, we can always answer probabilistic queries using standard Markov network inference methods on the instantiated network. We have extended and adapted several of these standard methods to take particular advantage of the logical structure in a Markov logic network, yielding tremendous savings in memory and time. We describe these algorithms in this section.

For many queries, only a small subset of the instantiated Markov network is relevant. In such cases, we need not instantiate or even consider the entire MLN. The proof DAG from a logical inference engine can be used to generate the set of ground formulas and atoms relevant to a particular query. Together with the MLN weights, this can be used to generate a sub-network to answer the probabilistic query. In this way, Markov logic can easily be paired with traditional logical inference methods. This method, traditionally known as knowledge-based model construction (KBMC) [27], allows us to potentially reason efficiently over a very large knowledge base (like the Semantic Web), as long as only a small fraction of it is relevant to the query. In our descriptions, we will assume that inference is done over the entire MLN, but our methods apply to the sub-network case as well.

5.1 MAP/MPE Inference

In the remainder of this chapter, we assume that the MLN is in function-free clausal form for convenience, but these methods can be applied to other MLNs as well. A basic inference task is finding the most probable state of the world given some evidence. (This is known as MAP inference in the Markov network literature, and MPE inference in the Bayesian network literature.) Because of the form of Equation 3, in Markov logic this reduces to finding the truth assignment that maximizes the sum of weights of satisfied clauses. This can be done using any weighted satisfiability solver, and (remarkably) need not be more expensive than standard logical inference by model checking. (In fact, it can be faster, if some hard constraints are softened.) We have successfully used MaxWalkSAT, a weighted variant of the WalkSAT local-search satisfiability solver, which can solve hard problems with hundreds of thousands of variables in minutes [13]. MaxWalkSAT performs this stochastic search by picking an unsatisfied clause at random and flipping the truth value of one of the atoms in it. With a certain probability, the atom is chosen randomly; otherwise, the atom is chosen to maximize the sum of satisfied clause weights when flipped. This combination of random and greedy steps allows MaxWalkSAT to avoid getting stuck in local optima while searching. Pseudocode for MaxWalkSAT is shown in Algorithm 1. $\Delta\text{Cost}(v)$ computes the change in the sum of weights of unsatisfied clauses that results from flipping variable v in the current solution. $\text{Uniform}(0,1)$ returns a uniform deviate from the interval $[0, 1]$.

One problem with this approach is that it requires propositionalizing the domain (i.e., grounding all atoms and clauses in all possible ways), which consumes memory exponential in the arity of the clauses. We have overcome this by developing LazySAT, a lazy version of MaxWalkSAT which grounds atoms and clauses only as needed [42]. This takes advantage of the sparseness of relational domains, where most atoms are false and most clauses are trivially satisfied. For example, in the domain of scientific research, most groundings of the atom $\text{Author}(\text{person}, \text{paper})$ are false, and most groundings of the clause $\text{Author}(\text{person1}, \text{paper}) \wedge \text{Author}(\text{person2}, \text{paper}) \Rightarrow \text{Coauthor}(\text{person1}, \text{person2})$ are satisfied. In LazySAT, the memory cost does not scale with the number of possible clause groundings, but only with the number of groundings that are potentially unsatisfied at some point in the search.

Algorithm 2 gives pseudo-code for LazySAT, highlighting the places where it differs from MaxWalkSAT. LazySAT maintains a set of *active atoms* and a set of *active clauses*. A clause is active if it can be made unsatisfied by flipping zero or more of its active atoms. (Thus, by definition, an unsatisfied clause is always active.) An atom is active if it is in the initial set of active atoms, or if it was flipped at some point in the search. The initial active atoms are all those appearing in clauses that are unsatisfied if only the atoms in the database are true, and all others are false. The unsatisfied clauses are obtained by simply going through each possible grounding of all the first-order clauses and materializing the groundings that are unsatisfied; search is pruned as soon as the partial grounding of a clause is satisfied. Given the initial active atoms, the definition

Algorithm 1 MaxWalkSAT(*weighted_clauses*, *max_flips*, *max_tries*, *target*, *p*)

```
vars ← variables in weighted_clauses
for i ← 1 to max_tries do
  soln ← a random truth assignment to vars
  cost ← sum of weights of unsatisfied clauses in soln
  for i ← 1 to max_flips do
    if cost ≤ target then
      return “Success, solution is”, soln
    end if
    c ← a randomly chosen unsatisfied clause
    if Uniform(0,1) < p then
      vf ← a randomly chosen variable from c
    else
      for each variable v in c do
        compute DeltaCost(v)
      end for
      vf ← v with lowest DeltaCost(v)
    end if
    soln ← soln with vf flipped
    cost ← cost + DeltaCost(vf)
  end for
end for
return “Failure, best assignment is”, best soln found
```

of active clause requires that some clauses become active, and these are found using a similar process (with the difference that, instead of checking whether a ground clause is unsatisfied, we check whether it should be active). Each run of LazySAT is initialized by assigning random truth values to the active atoms. This differs from MaxWalkSAT, which assigns random values to all atoms. However, the LazySAT initialization is a valid MaxWalkSAT initialization, and we have verified experimentally that the two give very similar results. Given the same initialization, the two algorithms will produce exactly the same results.

At each step in the search, the variable that is flipped is activated, as are any clauses that by definition should become active as a result. When evaluating the effect on cost of flipping a variable *v*, if *v* is active then all of the relevant clauses are already active, and DeltaCost(*v*) can be computed as in MaxWalkSAT. If *v* is inactive, DeltaCost(*v*) needs to be computed using the knowledge base. This is done by retrieving from the KB all first-order clauses containing the atom that *v* is a grounding of, and grounding each such clause with the constants in *v* and all possible groundings of the remaining variables. As before, we prune search as soon as a partial grounding is satisfied, and add the appropriate multiple of the clause weight to DeltaCost(*v*). (A similar process is used to activate clauses.) While this process is costlier than using pre-grounded clauses, it is amortized over many tests of active variables. In typical satisfiability problems, a small core of “problem” clauses is repeatedly tested, and when this is the case LazySAT will be quite efficient.

Algorithm 2 LazySAT(*weighted_KB*, *DB*, *max_flips*, *max_tries*, *target*, *p*)

```
for  $i \leftarrow 1$  to  $max\_tries$  do
   $active\_atoms \leftarrow$  atoms in clauses not satisfied by  $DB$ 
   $active\_clauses \leftarrow$  clauses activated by  $active\_atoms$ 
   $soln \leftarrow$  a random truth assignment to  $active\_atoms$ 
   $cost \leftarrow$  sum of weights of unsatisfied clauses in  $soln$ 
  for  $i \leftarrow 1$  to  $max\_flips$  do
    if  $cost \leq target$  then
      return "Success, solution is",  $soln$ 
    end if
     $c \leftarrow$  a randomly chosen unsatisfied clause
    if Uniform(0,1) <  $p$  then
       $v_f \leftarrow$  a randomly chosen variable from  $c$ 
    else
      for each variable  $v$  in  $c$  do
        compute  $\Delta Cost(v)$ , using  $weighted\_KB$  if  $v \notin active\_atoms$ 
      end for
       $v_f \leftarrow v$  with lowest  $\Delta Cost(v)$ 
    end if
    if  $v_f \notin active\_atoms$  then
      add  $v_f$  to  $active\_atoms$ 
      add clauses activated by  $v_f$  to  $active\_clauses$ 
    end if
     $soln \leftarrow soln$  with  $v_f$  flipped
     $cost \leftarrow cost + \Delta Cost(v_f)$ 
  end for
end for
return "Failure, best assignment is", best  $soln$  found
```

At each step, LazySAT flips the same variable that MaxWalkSAT would, and hence the result of the search is the same. The memory cost of LazySAT is on the order of the maximum number of clauses active at the end of a run of flips. (The memory required to store the active atoms is dominated by the memory required to store the active clauses, since each active atom appears in at least one active clause.)

Experiments on entity resolution and planning problems show that this can yield very large memory reductions, and these reductions increase with domain size [42]. For domains whose full instantiations fit in memory, running time is comparable; as problems become larger, full instantiation for MaxWalkSAT becomes impossible.

5.2 Marginal and Conditional Probabilities

Another key inference task is computing the probability that a formula holds, given an MLN and set of constants, and possibly other formulas as evidence.

By definition, the probability of a formula is the sum of the probabilities of the worlds where it holds, and computing it by brute force requires time exponential in the number of possible ground atoms. An approximate but more efficient alternative is to use Markov chain Monte Carlo (MCMC) inference [11], which samples a sequence of states according to their probabilities, and counting the fraction of sampled states where the formula holds. This can be extended to conditioning on other formulas by rejecting any state that violates one of them.

For the remainder of the chapter, we focus on the typical case where the evidence is a conjunction of ground atoms. In this scenario, further efficiency can be gained by applying a generalization of knowledge-based model construction [49]. This constructs only the minimal subset of the ground network required to answer the query, and runs MCMC (or any other probabilistic inference method) on it. The network is constructed by checking if the atoms that the query formula directly depends on are in the evidence. If they are, the construction is complete. Those that are not are added to the network, and we in turn check the atoms they depend on. This process is repeated until all relevant atoms have been retrieved. While in the worst case it yields no savings, in practice it can vastly reduce the time and memory required for inference. See Richardson and Domingos [36] for details.

One problem with applying MCMC to MLNs is that it breaks down in the presence of deterministic or near-deterministic dependencies (as do other probabilistic inference methods, e.g., belief propagation [52]). Deterministic dependencies break up the space of possible worlds into regions that are not reachable from each other, violating a basic requirement of MCMC. Near-deterministic dependencies greatly slow down inference, by creating regions of low probability that are very difficult to traverse. Running multiple chains with random starting points does not solve this problem, because it does not guarantee that different regions will be sampled with frequency proportional to their probability, and there may be a very large number of regions.

We have successfully addressed this problem by combining MCMC with satisfiability testing in the MC-SAT algorithm [32]. MC-SAT is a *slice sampling* MCMC algorithm. It uses a combination of satisfiability testing and simulated annealing to sample from the slice. The advantage of using a satisfiability solver (WalkSAT) is that it efficiently finds isolated modes in the distribution, and as a result the Markov chain mixes very rapidly. The slice sampling scheme ensures that detailed balance is (approximately) preserved.

MC-SAT is orders of magnitude faster than standard MCMC methods such as Gibbs sampling and simulated tempering, and is applicable to any model that can be expressed in Markov logic, including many standard models in statistical physics, vision, natural language processing, social network analysis, spatial statistics, etc.

Slice sampling [5] is an instance of a widely used approach in MCMC inference that introduces *auxiliary variables* to capture the dependencies between observed variables. For example, to sample from $P(X = x) = (1/Z) \prod_k \phi_k(x_{\{k\}})$, we can define $P(X = x, U = u) = (1/Z) \prod_k I_{[0, \phi_k(x_{\{k\}})]}(u_k)$, where ϕ_k is the k th

Algorithm 3 MC-SAT(*clauses, weights, num_samples*)

```
 $x^{(0)} \leftarrow \text{Satisfy}(\text{hard clauses})$   
for  $i \leftarrow 1$  to  $\text{num\_samples}$  do  
   $M \leftarrow \emptyset$   
  for all  $c_k \in \text{clauses}$  satisfied by  $x^{(i-1)}$  do  
    With probability  $1 - e^{-w_k}$  add  $c_k$  to  $M$   
  end for  
  Sample  $x^{(i)} \sim \mathcal{U}_{\text{SAT}(M)}$   
end for
```

potential function, u_k is the k th auxiliary variable, $I_{[a,b]}(u_k) = 1$ if $a \leq u_k \leq b$, and $I_{[a,b]}(u_k) = 0$ otherwise. The marginal distribution of X under this joint is $P(X=x)$, so to sample from the original distribution it suffices to sample from $P(x, u)$ and ignore the u values. $P(u_k|x)$ is uniform in $[0, \phi_k(x_{\{k\}})]$, and thus easy to sample from. The main challenge is to sample x given u , which is uniform among all \mathcal{X} that satisfies $\phi_k(x_{\{k\}}) \geq u_k$ for all k . MC-SAT uses SampleSAT [48] to do this. In each sampling step, MC-SAT takes the set of all ground clauses satisfied by the current state of the world and constructs a subset, M , that must be satisfied by the next sampled state of the world. (For the moment we will assume that all clauses have positive weight.) Specifically, a satisfied ground clause is included in M with probability $1 - e^{-w}$, where w is the clause’s weight. We then take as the next state a uniform sample from the set of states $\text{SAT}(M)$ that satisfy M . (Notice that $\text{SAT}(M)$ is never empty, because it always contains at least the current state.) Algorithm 3 gives pseudo-code for MC-SAT. \mathcal{U}_S is the uniform distribution over set S . At each step, all hard clauses are selected with probability 1, and thus all sampled states satisfy them. Negative weights are handled by noting that a clause with weight $w < 0$ is equivalent to its negation with weight $-w$, and a clause’s negation is the conjunction of the negations of all of its literals. Thus, instead of checking whether the clause is satisfied, we check whether its negation is satisfied; if it is, with probability $1 - e^w$ we select all of its negated literals, and with probability e^w we select none.

It can be shown that MC-SAT satisfies the MCMC criteria of detailed balance and ergodicity [32], assuming a perfect uniform sampler. In general, uniform sampling is #P-hard and SampleSAT [48] only yields approximately uniform samples. However, experiments show that MC-SAT is still able to produce very accurate probability estimates, and its performance is not very sensitive to the parameter setting of SampleSAT.

We have applied the ideas of LazySAT to implement a lazy version of MC-SAT that avoids grounding unnecessary atoms and clauses. A working version of this algorithm is present in the open-source Alchemy system [16].

It is also possible to carry out lifted first-order probabilistic inference (akin to resolution) in Markov logic [3]. These methods speed up inference by reasoning at the first-order level about groups of indistinguishable objects rather than propositionalizing the entire domain. This is particularly applicable when the population size is given but little is known about most individual members.

6 Learning

In this section, we discuss methods for automatically learning weights, refining formulas, and constructing new formulas from data. Of course, learning is but one method for generating an MLN. In a distributed knowledge base such as the Semantic Web, formulas could come from many different sources and their weights could be set by the sources themselves or using credibility or trust propagation (e.g., [35]). When data is available, learning methods allow us to automatically adjust weights and refine or add formulas to an MLN.

6.1 Generative Weight Learning

MLN weights can be learned generatively by maximizing the likelihood of a relational database (Equation 3). This relational database consists of one or more “possible worlds” that form our training examples. Note that we can learn to generalize from even a single example because the clause weights are shared across their many respective groundings. This is essential when the training data is a single network, such as in the Semantic Web. The gradient of the log-likelihood with respect to the weights is

$$\frac{\partial}{\partial w_i} \log P_w(X=x) = n_i(x) - \sum_{x'} P_w(X=x') n_i(x') \quad (4)$$

where the sum is over all possible databases x' , and $P_w(X=x')$ is $P(X=x')$ computed using the current weight vector $w = (w_1, \dots, w_i, \dots)$. In other words, the i th component of the gradient is simply the difference between the number of true groundings of the i th formula in the data and its expectation according to the current model. In the generative case, even approximating these expectations tends to be prohibitively expensive or inaccurate due to the large state space. Instead, we maximize the pseudo-likelihood of the data, a widely-used alternative [2]. If x is a possible world (relational database) and x_l is the l th ground atom’s truth value, the pseudo-log-likelihood of x given weights w is

$$\log P_w^*(X=x) = \sum_{l=1}^n \log P_w(X_l=x_l | MB_x(X_l)) \quad (5)$$

where $MB_x(X_l)$ is the state of X_l ’s Markov blanket in the data (i.e., the truth values of the ground atoms it appears in some ground formula with). Computing the pseudo-likelihood and its gradient does not require inference, and is therefore much faster. Combined with the L-BFGS optimizer [20], pseudo-likelihood yields efficient learning of MLN weights even in domains with millions of ground atoms [36]. However, the pseudo-likelihood parameters may lead to poor results when long chains of inference are required.

In order to reduce overfitting, we penalize each weight with a Gaussian prior. We apply this strategy not only to generative learning, but to all of our weight learning methods, even those embedded within structure learning.

6.2 Discriminative Weight Learning

Discriminative learning is an attractive alternative to pseudo-likelihood. In many applications, we know *a priori* which atoms will be evidence and which ones will be queried, and the goal is to correctly predict the latter given the former. If we partition the ground atoms in the domain into a set of evidence atoms X and a set of query atoms Y , the *conditional likelihood (CLL)* of Y given X is $P(y|x) = (1/Z_x) \exp(\sum_{i \in F_Y} w_i n_i(x, y)) = (1/Z_x) \exp(\sum_{j \in G_Y} w_j g_j(x, y))$, where F_Y is the set of all MLN clauses with at least one grounding involving a query atom, $n_i(x, y)$ is the number of true groundings of the i th clause involving query atoms, G_Y is the set of ground clauses in $M_{L,C}$ involving query atoms, and $g_j(x, y) = 1$ if the j th ground clause is true in the data and 0 otherwise. The gradient of the CLL is

$$\begin{aligned} \frac{\partial}{\partial w_i} \log P_w(y|x) &= n_i(x, y) - \sum_{y'} P_w(y'|x) n_i(x, y') \\ &= n_i(x, y) - E_w[n_i(x, y)] \end{aligned} \tag{6}$$

In the conditional case, we can approximate the expected counts $E_w[n_i(x, y)]$ using either the MAP state (i.e., the most probable state of y given x) or by averaging over several MC-SAT samples. The MAP approximation is inspired by the voted perceptron algorithm proposed by Collins [4] for discriminatively learning hidden Markov models. We can apply a similar algorithm to MLNs using MaxWalkSAT to find the approximate MAP state, following the approximate gradient for a fixed number of iterations, and averaging the weights across all iterations to combat overfitting [40]. We get the best results, however, by applying a version of the scaled conjugate gradient algorithm [26]. We use a small number of MC-SAT samples to approximate the gradient and Hessian matrix, and use the inverse diagonal hessian as a preconditioner. See Lowd and Domingos [22] for more details and results.

6.3 Structure Learning

The structure of a Markov logic network is the set of formulas or clauses to which we attach weights. While this knowledge base is often specified by one or more experts, such knowledge is not always accurate or complete. In addition to learning weights for the provided clauses, we can revise or extend the MLN structure with new clauses learned from data. The inductive logic programming (ILP) community has developed many methods for learning logical rules from data. However, since an MLN represents a probability distribution, much better results are obtained by using an evaluation function based on pseudo-likelihood, rather than typical ILP ones like accuracy and coverage [14]. Log-likelihood or conditional log-likelihood are potentially better evaluation functions, but are vastly more expensive to compute. In experiments on two real-world datasets, our MLN structure learning algorithm found better MLN rules than the standard

ILP algorithms CLAUDIEN [6], FOIL [34], and Aleph [45], and even a hand-written knowledge base.

MLN structure learning can start from an empty network or from an existing KB. Either way, we have found it useful to start by adding all unit clauses (single atoms) to the MLN. The weights of these capture (roughly speaking) the marginal distributions of the atoms, allowing the longer clauses to focus on modeling atom dependencies. To extend this initial model, we either repeatedly find the best clause using beam search and add it to the MLN, or add all “good” clauses of length l before trying clauses of length $l + 1$. Candidate clauses are formed by adding each predicate (negated or otherwise) to each current clause, with all possible combinations of variables, subject to the constraint that at least one variable in the new predicate must appear in the current clause. Hand-coded clauses are also modified by removing predicates.

Recently, Mihalkova and Mooney [25] introduced BUSL, an alternative, bottom-up structure learning algorithm for Markov logic. Instead of blindly constructing candidate clauses one literal at a time, they let the training data guide and constrain clause construction. First, they use a propositional Markov network structure learner to generate a graph of relationships among atoms. Then they generate clauses from paths in this graph. In this way, BUSL focuses on clauses that have support in the training data. In experiments on three datasets, BUSL evaluated many fewer candidate clauses than our top-down algorithm, ran more quickly, and learned more accurate models.

We are currently investigating further approaches to learning MLNs, including automatically inventing new predicates (or, in statistical terms, discovering hidden variables) [15].

7 Applications

We have already applied Markov logic to a variety of problems relevant to the Semantic Web, including link prediction and collective classification, for filling in missing attributes and relationships; entity resolution, for matching equivalent entities that have different names; information extraction, for adding structure to raw or semi-structured text; and other problems [36, 40, 14, 41, 32, 33]. Even our simple Friends and Smokers example touches on link prediction, collective classification, and social network analysis. In this section, we will show in detail how Markov logic can be used to build state-of-the-art models for entity resolution and information extraction, and present experimental results on real-world citation data.

Others have also applied Markov logic in a variety of areas. A system based on it recently won a competition on information extraction for biology [37]. Cycorp has used it to make parts of the Cyc knowledge base probabilistic [24]. The CALO project is using it to integrate probabilistic learning and inference across many components [9]. Of particular relevance to the Semantic Web is the recent work of Wu and Weld [51] on automatically refining the Wikipedia infobox ontology.

7.1 Entity Resolution

The application to entity resolution illustrates well the power of Markov logic [41]. Entity resolution is the problem of determining which observations (e.g., database records, noun phrases, video regions, etc.) correspond to the same real-world objects. This is an important and difficult task even on small, well-defined, and well-maintained databases. In the Semantic Web, automatically determining which objects, fields, and types are equivalent becomes much harder since the data may come from many different sources with varied quality. Manual annotation does not scale, so automatically determining these relationships is essential for maintaining connectedness in the Semantic Web.

Entity resolution is typically done by forming a vector of properties for each pair of observations, using a learned classifier (such as logistic regression) to predict whether they match, and applying transitive closure. Markov logic yields an improved solution simply by applying the standard logical approach of removing the unique names assumption and introducing the equality predicate and its axioms: equality is reflexive, symmetric and transitive; groundings of a predicate with equal constants have the same truth values; and constants appearing in a ground predicate with equal constants are equal. This last axiom is not valid in logic, but captures a useful statistical tendency. For example, if two papers are the same, their authors are the same; and if two authors are the same, papers by them are more likely to be the same. Weights for different instances of these axioms can be learned from data. Inference over the resulting MLN, with entity properties and relations as the evidence and equality atoms as the query, naturally combines logistic regression and transitive closure. Most importantly, it performs *collective* entity resolution, where resolving one pair of entities helps to resolve pairs of related entities.

As a concrete example, consider the task of deduplicating a citation database in which each citation has author, title, and venue fields. We can represent the domain structure with eight relations: `Author(bib, author)`, `Title(bib, title)`, and `Venue(bib, venue)` relate citations to their fields; `HasWord(author/title/venue, word)` indicates which words are present in each field; `SameAuthor(author, author)`, `SameTitle(title, title)`, and `SameVenue(venue, venue)` represent field equivalence; and `SameBib(bib, bib)` represents citation equivalence. The truth values of all relations except for the equivalence relations are provided as background theory. The objective is to predict the `SameBib` relation.

We begin with a logistic regression model to predict citation equivalence based on the words in the fields. This is easily expressed in Markov logic by rules such as the following:

$$\begin{aligned} & \text{Title}(b1, t1) \wedge \text{Title}(b2, t2) \wedge \text{HasWord}(t1, +\text{word}) \\ & \wedge \text{HasWord}(t2, +\text{word}) \Rightarrow \text{SameBib}(b1, b2) \end{aligned}$$

The ‘+’ operator here generates a separate rule (and with it, a separate learnable weight) for each constant of the appropriate type. When given a positive weight, each of these rules increases the probability that two citations with a particular title word in common are equivalent. We can construct similar rules for other

fields. Note that we may learn negative weights for some of these rules, just as logistic regression may learn negative feature weights. Transitive closure consists of a single rule:

$$\text{SameBib}(b1, b2) \wedge \text{SameBib}(b2, b3) \Rightarrow \text{SameBib}(b1, b3)$$

This model is similar to the standard solution, but has the advantage that the classifier is learned in the context of the transitive closure operation.

We can construct similar rules to predict the equivalence of two fields as well. The usefulness of Markov logic is shown further when we link field equivalence to citation equivalence:

$$\begin{aligned} \text{Author}(b1, a1) \wedge \text{Author}(b2, a2) \wedge \text{SameBib}(b1, b2) &\Rightarrow \text{SameAuthor}(a1, a2) \\ \text{Author}(b1, a1) \wedge \text{Author}(b2, a2) \wedge \text{SameAuthor}(a1, a2) &\Rightarrow \text{SameBib}(b1, b2) \end{aligned}$$

The above rules state that if two citations are the same, their authors should be the same, and that citations with the same author are more likely to be the same. The last rule is not valid in logic, but captures a useful statistical tendency.

Most importantly, the resulting model can now perform *collective* entity resolution, where resolving one pair of entities helps to resolve pairs of related entities. For example, inferring that a pair of citations are equivalent can provide evidence that the names *AAAI-06* and *21st Natl. Conf. on AI* refer to the same venue, even though they are superficially very different. This equivalence can then aid in resolving other entities.

Experiments on citation databases like Cora and BibServ.org show that these methods can greatly improve accuracy, particularly for entity types that are difficult to resolve in isolation as in the above example [41]. Due to the large number of words and the high arity of the transitive closure formula, these models have thousands of weights and ground millions of clauses during learning, even after using canopies to limit the number of comparisons considered. Learning at this scale is still reasonably efficient: preconditioned scaled conjugate gradient with MC-SAT for inference converges within a few hours [22].

7.2 Information Extraction

In this citation example, it was assumed that the fields were manually segmented in advance. The goal of information extraction is to extract database records starting from raw text or semi-structured data sources. This has many applications for the Semantic Web, including using the vast amount of unstructured information on the Web to bootstrap the Semantic Web. Information extraction could also be used to segment labeled fields, such as “name,” into more specific fields, such as “first name,” “last name,” and “title.”

Traditionally, information extraction proceeds by first segmenting each candidate record separately, and then merging records that refer to the same entities. Such a pipeline architecture is adopted by many AI systems in natural language processing, speech recognition, vision, robotics, etc. Markov logic allows us to perform the two tasks jointly [33]. This enables us to use the segmentation of

one candidate record to help segment similar ones. For example, resolving a well-segmented field with a less-clear one can disambiguate the latter’s boundaries. We will continue with the example of citations, but similar ideas could be applied to other data sources, such as Web pages or emails.

The main evidence predicate in the information extraction MLN is $\text{Token}(\mathbf{t}, \mathbf{i}, \mathbf{c})$, which is true iff token \mathbf{t} appears in the \mathbf{i} th position of the \mathbf{c} th citation. A token can be a word, date, number, etc. Punctuation marks are not treated as separate tokens; rather, the predicate $\text{HasPunc}(\mathbf{c}, \mathbf{i})$ is true iff a punctuation mark appears immediately after the \mathbf{i} th position in the \mathbf{c} th citation. The query predicates are $\text{InField}(\mathbf{i}, \mathbf{f}, \mathbf{c})$ and $\text{SameCitation}(\mathbf{c}, \mathbf{c}')$. $\text{InField}(\mathbf{i}, \mathbf{f}, \mathbf{c})$ is true iff the \mathbf{i} th position of the \mathbf{c} th citation is part of field \mathbf{f} , where $\mathbf{f} \in \{\text{Title}, \text{Author}, \text{Venue}\}$, and inferring it performs segmentation. $\text{SameCitation}(\mathbf{c}, \mathbf{c}')$ is true iff citations \mathbf{c} and \mathbf{c}' represent the same publication, and inferring it performs entity resolution.

Our segmentation model is essentially a hidden Markov model (HMM) with enhanced ability to detect field boundaries. The observation matrix of the HMM correlates tokens with fields, and is represented by the simple rule

$$\text{Token}(+\mathbf{t}, \mathbf{i}, \mathbf{c}) \Rightarrow \text{InField}(\mathbf{i}, +\mathbf{f}, \mathbf{c})$$

If this rule was learned in isolation, the weight of the (t, f) th instance would be $\log(p_{tf}/(1-p_{tf}))$, where p_{tf} is the corresponding entry in the HMM observation matrix. In general, the transition matrix of the HMM is represented by a rule of the form

$$\text{InField}(\mathbf{i}, +\mathbf{f}, \mathbf{c}) \Rightarrow \text{InField}(\mathbf{i} + 1, +\mathbf{f}', \mathbf{c})$$

However, we (and others, e.g., [12]) have found that for segmentation it suffices to capture the basic regularity that consecutive positions tend to be part of the same field. Thus we replace \mathbf{f}' by \mathbf{f} in the formula above. We also impose the condition that a position in a citation string can be part of at most one field; it may be part of none.

The main shortcoming of this model is that it has difficulty pinpointing field boundaries. Detecting these is key for information extraction, and a number of approaches use rules designed specifically for this purpose (e.g., [17]). In citation matching, boundaries are usually marked by punctuation symbols. This can be incorporated into the MLN by modifying the rule above to

$$\text{InField}(\mathbf{i}, +\mathbf{f}, \mathbf{c}) \wedge \neg \text{HasPunc}(\mathbf{c}, \mathbf{i}) \Rightarrow \text{InField}(\mathbf{i} + 1, +\mathbf{f}, \mathbf{c})$$

The $\neg \text{HasPunc}(\mathbf{c}, \mathbf{i})$ precondition prevents propagation of fields across punctuation marks. Because propagation can occur differentially to the left and right, the MLN also contains the reverse form of the rule. In addition, to account for commas being weaker separators than other punctuation, the MLN includes versions of these rules with $\text{HasComma}()$ instead of $\text{HasPunc}()$.

Finally, the MLN contains rules capturing a variety of knowledge about citations: the first two positions of a citation are usually in the author field, and the middle one in the title; initials (e.g., “J.”) tend to appear in either the author or the venue field; positions preceding the last non-venue initial are usually

Table 2. CiteSeer entity resolution: cluster recall on each section.

Approach	Constr.	Face	Reason.	Reinfor.
Fellegi-Sunter	84.3	81.4	71.3	50.6
Lawrence et al. (1999)	89	94	86	79
Pasula et al. (2002)	93	97	96	94
Wellner et al. (2004)	95.1	96.9	93.7	94.7
Joint MLN	96.0	97.1	95.1	96.7

not part of the title or venue; and positions after the first venue keyword (e.g., “Proceedings”, “Journal”) are usually not part of the author or title.

By combining this segmentation model with our entity resolution model from before, we can exploit relational information as part of the segmentation process. In practice, something a little more sophisticated is necessary to get good results on real data. In Poon and Domingos [33], we define predicates and rules specifically for passing information between the stages, as opposed to just using the existing `InField()` outputs. This leads to a “higher bandwidth” of communication between segmentation and entity resolution, without letting excessive segmentation noise through. We also define an additional predicate and modify rules to better exploit information from similar citations during the segmentation process. See [33] for further details.

We evaluated this model on the CiteSeer and Cora datasets. For entity resolution in CiteSeer, we measured *cluster recall* for comparison with previously published results. Cluster recall is the fraction of clusters that are correctly output by the system after taking transitive closure from pairwise decisions. For entity resolution in Cora, we measured both cluster recall and pairwise recall/precision. In both datasets we also compared with a “standard” Fellegi-Sunter model (see [41]), learned using logistic regression, and with oracle segmentation as the input.

In both datasets, joint inference improved accuracy and our approach outperformed previous ones. Table 2 shows that our approach outperforms previous ones on CiteSeer entity resolution. (Results for Lawrence et al. (1999) [19], Pasula et al. (2002) [29] and Wellner et al. (2004) [50] are taken from the corresponding papers.) This is particularly notable given that the models of [29] and [50] involved considerably more knowledge engineering than ours, contained more learnable parameters, and used additional training data.

Table 3 shows that our entity resolution approach easily outperforms Fellegi-Sunter on Cora, and has very high pairwise recall/precision.

8 The Alchemy System

The inference and learning algorithms described in the previous sections are publicly available in the open-source Alchemy system [16]. Alchemy makes it possible to define sophisticated probabilistic models with a few formulas, and

Table 3. Cora entity resolution: pairwise recall/precision and cluster recall.

Approach	Pairwise Rec./Prec.	Cluster Recall
Fellegi-Sunter	78.0 / 97.7	62.7
Joint MLN	94.3 / 97.0	78.1

to add probability to a first-order knowledge base by learning weights from a relevant database. It can also be used for purely logical or purely statistical applications, and for teaching AI. From the user’s point of view, *Alchemy* provides a full spectrum of AI tools in an easy-to-use, coherent form. From the researcher’s point of view, *Alchemy* makes it possible to easily integrate a new inference or learning algorithm, logical or statistical, with a full complement of other algorithms that support it or make use of it.

Alchemy can be viewed as a declarative programming language akin to *Prolog*, but with a number of key differences: the underlying inference mechanism is model checking instead of theorem proving; the full syntax of first-order logic is allowed, rather than just Horn clauses; and, most importantly, the ability to handle uncertainty and learn from data is already built in. Table 4 compares *Alchemy* with *Prolog* and *BUGS* [23], one of the most popular toolkits for Bayesian modeling and inference.

Table 4. A comparison of *Alchemy*, *Prolog* and *BUGS*.

Aspect	<i>Alchemy</i>	<i>Prolog</i>	<i>BUGS</i>
Representation	First-order logic + Markov nets	Horn clauses	Bayes nets
Inference	Model checking, MCMC	Theorem proving	MCMC
Learning	Parameters and structure	No	Parameters
Uncertainty	Yes	No	Yes
Relational	Yes	Yes	No

9 Current and Future Research Directions

We are actively researching better learning and inference methods for Markov logic, as well as extensions of the representation that increase its generality and power.

Exact methods for learning and inference are usually intractable in Markov logic, but we would like to see better, more efficient approximations along with the automatic application of exact methods when feasible.

One method of particular interest is lifted inference. In short, we would like to reason with clusters of nodes for which we have exactly the same amount

of information. The inspiration is from lifted resolution in first order logic, but must be extended to handle uncertainty. Prior work on lifted inference such as [31] and [3] mainly focused on exact inference which can be quite slow. We have recently extended loopy belief propagation, an approximate inference method for probabilistic graphical models, to perform lifted inference in Markov logic networks [44]. When the amount of evidence is limited, this can speed up inference by many orders of magnitude.

We are also working to develop a general framework for decision-making in relational domains. This can be accomplished in Markov logic by adding utility weights to formulas and finding the settings of all action predicates that jointly maximize expected utility. Decision-making is key to the original Semantic Web vision, which called for intelligent agents to act on the information they gathered.

Numerical attributes must be discretized to be used in Markov logic, but we have recently introduced methods to incorporate continuous random variables and features [46]. Continuous values could be useful in a variety of Semantic Web problems, such as incorporating numeric features into similarities for entity resolution, ontology alignment, or schema matching.

Current work also includes semi-supervised learning, and learning with incomplete data in general. The large amount of unlabeled data on the Web is an excellent resource that, properly exploited, could help bootstrap or enrich the Semantic Web.

10 Conclusion

The Semantic Web must deal with uncertainty from many sources, including inconsistent knowledge bases, incorrect or untrustworthy information, missing data, different ontologies and schemas, and more. Markov logic is a simple yet powerful approach for adding probability to logical representations such as those already used by the Semantic Web: Given a set of formulas, just add weights. We have developed a series of learning and inference algorithms for it, and successfully applied them in a number of domains. These algorithms are included in the open-source Alchemy system (available at alchemy.cs.washington.edu). We hope that Markov logic and its implementation in Alchemy will be of use to Semantic Web researchers and practitioners who wish to have the full spectrum of logical and statistical inference and learning techniques at their disposal, without having to develop every piece themselves.

11 Acknowledgements

This research was partly supported by DARPA grant FA8750-05-2-0283 (managed by AFRL), DARPA contract NBCH-D030010, NSF grant IIS-0534881, ONR grants N00014-02-1-0408 and N00014-05-1-0313, a Sloan Fellowship and NSF CAREER Award to the first author, and a Microsoft Research fellowship awarded to the second author. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily

representing the official policies, either expressed or implied, of DARPA, NSF, ONR, or the United States Government.

References

1. T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284 (5)(5):34–43, 2001.
2. J. Besag. Statistical analysis of non-lattice data. *The Statistician*, 24:179–195, 1975.
3. R. Braz, E. Amir, and D. Roth. Lifted first-order probabilistic inference. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 1319–1325, Edinburgh, UK, 2005. Morgan Kaufmann.
4. M. Collins. Discriminative training methods for hidden Markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, pages 1–8, Philadelphia, PA, 2002. ACL.
5. P. Damien, J. Wakefield, and S. Walker. Gibbs sampling for Bayesian non-conjugate and hierarchical models by auxiliary variables. *Journal of the Royal Statistical Society, Series B*, 61, 1999.
6. L. De Raedt and L. Dehaspe. Clausal discovery. *Machine Learning*, 26:99–146, 1997.
7. M. Dean, G. Schreiber, S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein. OWL web ontology language reference. 2004. <http://www.w3.org/TR/owl-ref/>.
8. S. Della Pietra, V. Della Pietra, and J. Lafferty. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:380–392, 1997.
9. T. Dietterich and X. Bao. Integrating multiple learning components through Markov logic. In *Proceedings of the Twenty-Third National Conference on Artificial Intelligence*, Chicago, IL, 2008. AAAI Press.
10. M. R. Genesereth and N. J. Nilsson. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, San Mateo, CA, 1987.
11. W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London, UK, 1996.
12. T. Grenager, D. Klein, and C. D. Manning. Unsupervised learning of field segmentation models for information extraction. In *Proceedings of the Forty-Third Annual Meeting on Association for Computational Linguistics*, pages 371–378, Ann Arbor, Michigan, 2005. Association for Computational Linguistics.
13. H. Kautz, B. Selman, and Y. Jiang. A general stochastic approach to solving problems with hard and soft constraints. In D. Gu, J. Du, and P. Pardalos, editors, *The Satisfiability Problem: Theory and Applications*, pages 573–586. American Mathematical Society, New York, NY, 1997.
14. S. Kok and P. Domingos. Learning the structure of Markov logic networks. In *Proceedings of the Twenty-Second International Conference on Machine Learning*, pages 441–448, Bonn, Germany, 2005. ACM Press.
15. S. Kok and P. Domingos. Statistical predicate invention. In *Proceedings of the Twenty-Fourth International Conference on Machine Learning*, pages 433–440, Corvallis, OR, 2007. ACM Press.

16. S. Kok, M. Sumner, M. Richardson, P. Singla, H. Poon, D. Lowd, and P. Domingos. The Alchemy system for statistical relational AI. Technical report, Department of Computer Science and Engineering, University of Washington, Seattle, WA, 2007. <http://alchemy.cs.washington.edu>.
17. N. Kushmerick. Wrapper induction: Efficiency and expressiveness. *Artificial Intelligence*, 118(1-2):15–68, 2000.
18. N. Lavrač and S. Džeroski. *Inductive Logic Programming: Techniques and Applications*. Ellis Horwood, Chichester, UK, 1994.
19. S. Lawrence, K. Bollacker, and C. L. Giles. Autonomous citation matching. In *Proceedings of the Third International Conference on Autonomous Agents*, New York, 1999. ACM Press.
20. D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(3):503–528, 1989.
21. J. W. Lloyd. *Foundations of Logic Programming*. Springer, Berlin, Germany, 1987.
22. D. Lowd and P. Domingos. Efficient weight learning for Markov logic networks. In *Proceedings of the Eleventh European Conference on Principles and Practice of Knowledge Discovery in Databases*, pages 200–211, Warsaw, Poland, 2007. Springer.
23. D. J. Lunn, A. Thomas, N. Best, and D. Spiegelhalter. WinBUGS – a Bayesian modeling framework: concepts, structure, and extensibility. *Statistics and Computing*, 10:325–337, 2000.
24. C. Matuszek and M. Witbrock. Personal communication. 2006.
25. L. Mihalkova and R. Mooney. Bottom-up learning of Markov logic network structure. In *Proceedings of the Twenty-Fourth International Conference on Machine Learning*, pages 625–632, Corvallis, OR, 2007. ACM Press.
26. M. Møller. A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, 6:525–533, 1993.
27. L. Ngo and P. Haddawy. Answering queries from context-sensitive probabilistic knowledge bases. *Theoretical Computer Science*, 171:147–177, 1997.
28. J. Nocedal and S. Wright. *Numerical Optimization*. Springer, New York, NY, 2006.
29. H. Pasula, B. Marthi, B. Milch, S. Russell, and I. Shpitser. Identity uncertainty and citation matching. In *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2002. MIT Press.
30. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco, CA, 1988.
31. D. Poole. First-order probabilistic inference. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 985–991, Acapulco, Mexico, 2003. Morgan Kaufmann.
32. H. Poon and P. Domingos. Sound and efficient inference with probabilistic and deterministic dependencies. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pages 458–463, Boston, MA, 2006. AAAI Press.
33. H. Poon and P. Domingos. Joint inference in information extraction. In *Proceedings of the Twenty-Second National Conference on Artificial Intelligence*, pages 913–918, Vancouver, Canada, 2007. AAAI Press.
34. J. R. Quinlan. Learning logical definitions from relations. *Machine Learning*, 5:239–266, 1990.
35. M. Richardson, R. Agrawal, and P. Domingos. Trust management for the Semantic Web. In *Proceedings of the Second International Semantic Web Conference*, pages 351–368, Sanibel Island, FL, 2003. Springer.
36. M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62:107–136, 2006.

37. S. Riedel and E. Klein. Genic interaction extraction with semantic and syntactic chains. In *Proceedings of the Fourth Workshop on Learning Language in Logic*, pages 69–74, Bonn, Germany, 2005. IMLS.
38. J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12:23–41, 1965.
39. D. Roth. On the hardness of approximate reasoning. *Artificial Intelligence*, 82:273–302, 1996.
40. P. Singla and P. Domingos. Discriminative training of Markov logic networks. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, pages 868–873, Pittsburgh, PA, 2005. AAAI Press.
41. P. Singla and P. Domingos. Entity resolution with Markov logic. In *Proceedings of the Sixth IEEE International Conference on Data Mining*, pages 572–582, Hong Kong, 2006. IEEE Computer Society Press.
42. P. Singla and P. Domingos. Memory-efficient inference in relational domains. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, Boston, MA, 2006. AAAI Press.
43. P. Singla and P. Domingos. Markov logic in infinite domains. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pages 368–375, Vancouver, Canada, 2007. AUAI Press.
44. P. Singla and P. Domingos. Lifted first-order belief propagation. In *Proceedings of the Twenty-Third National Conference on Artificial Intelligence*, Chicago, IL, 2008. AAAI Press.
45. A. Srinivasan. The Aleph manual. Technical report, Computing Laboratory, Oxford University, 2000.
46. J. Wang and P. Domingos. Hybrid Markov logic networks. In *Proceedings of the Twenty-Third National Conference on Artificial Intelligence*, Chicago, IL, 2008. AAAI Press. To appear.
47. S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK, 1994.
48. W. Wei, J. Erenrich, and B. Selman. Towards efficient sampling: Exploiting random walk strategies. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, San Jose, CA, 2004. AAAI Press.
49. M. Wellman, J. S. Breese, and R. P. Goldman. From knowledge bases to decision models. *Knowledge Engineering Review*, 7, 1992.
50. B. Wellner, A. McCallum, F. Peng, and M. Hay. An integrated, conditional model of information extraction and coreference with application to citation matching. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 593–601, Banff, Canada, 2004. AUAI Press.
51. F. Wu and D. Weld. Automatically refining the Wikipedia infobox ontology. In *17th International World Wide Web Conference (WWW-08)*, Beijing, China, 2008.
52. J. S. Yedidia, W. T. Freeman, and Y. Weiss. Generalized belief propagation. In T. Leen, T. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 689–695. MIT Press, Cambridge, MA, 2001.