

Structure Learning for Relational Logistic Regression: An Ensemble Approach

Nandini Ramanan

Indiana University, Bloomington

Gautam Kunapuli

The University of Texas at Dallas

Tushar Khot

Allen Institute for Artificial Intelligence

Bahare Fatemi

University of British Columbia

Seyed Mehran Kazemi

University of British Columbia

David Poole

University of British Columbia

Kristian Kersting

TU Darmstadt University

Sriraam Natarajan

The University of Texas at Dallas

Abstract

We consider the problem of learning Relational Logistic Regression (RLR). Unlike standard logistic regression, the features of RLRs are first-order formulae with associated weight vectors instead of scalar weights. We turn the problem of learning RLR to learning these vector-weighted formulae and develop a learning algorithm based on the recently successful functional-gradient boosting methods for probabilistic logic models. We derive the functional gradients and show how weights can be learned simultaneously in an efficient manner. Our empirical evaluation on standard and novel data sets demonstrates the superiority of our approach over other methods for learning RLR.

Introduction

Statistical Relational Learning models (SRL) (Getoor and Taskar 2007; Raedt et al. 2016) combine the representational power of logic with the ability of probability theory specifically, and statistical models in general to model noise and uncertainty. They have generally ranged from directed models (Kersting and De Raedt 2007; Koller 1999; Heckerman, Meek, and Koller 2007; Kazemi et al. 2014a; Neville and Jensen 2007; Kazemi and Poole 2018) to undirected models (Richardson and Domingos 2006; Taskar et al. 2007; Kimmig et al. 2012). We consider the more recent, well-understood directed model of **Relational Logistic Regression** (RLR) (Kazemi et al. 2014a; Kazemi et al. 2014b; Fatemi, Kazemi, and Poole 2016). One of the key advantages of RLR is that they scale well with population size unlike other methods such as Markov Logic Networks (Poole et al. 2014) and hence can potentially be used as a powerful modeling tool for many tasks.

While the models are attractive from the modeling perspective, learning these models is computationally intensive. This is due to the fact that (like the field of Inductive Logic Programming) learning occurs at multiple levels of abstraction, that of the level of objects, sub-group of objects and relations and possibly at the individual instances of the objects. Hence, most methods for learning these models have so far focused on the task of learning the so-called parameters (weights of the logistic function) where the rules (or

relational features) are provided by the human expert and the data is merely used to learn the parameters.

We consider the problem of full model learning, also known as **structure learning** of RLR models. A simple solution to learning these models could be to learn the rules separately using a logic learner and then employ the parameter learning strategies (Huynh and Mooney 2008). While reasonably easy to implement, the key issue is that the disconnect between rule and parameter learning can result in poor predictive performance as shown repeatedly in the literature (Natarajan et al. 2012; Richardson and Domingos 2006). Inspired by the success of non-parametric learning methods for SRL models, we develop a learning method for full model learning of RLR models.

More specifically, we develop a gradient-boosting technique for learning RLR models. We derive the gradients for the different weights of RLR and show how the rules of the logistic function are learned simultaneously with their corresponding weights. Unlike the standard adaptations of the functional gradients, RLR requires learning a different set of weights per rule in each gradient step and hence requires learning multiple weights jointly for a single rule. As we explain later, the gradients correspond to a set of vector weighted clauses that are learned in a sequential manner. We derive the gradients for these clauses and illustrate how to optimize them.

Each clause can be seen as a **relational feature** for the logistic function. We also note that RLR can be viewed as a probabilistic combination function in that it can stochastically combine the distributions due to different set of parents (in graphical model terminology). Hence, if our learning technique is employed in the context of learning joint models, our work can be seen as a new interpretation of learning boosted Relational Dependency Networks (RDNs) (Neville and Jensen 2007; Natarajan et al. 2012), where the standard aggregators are replaced with a logistic regression combination function which could potentially yield interesting insights into directed SRL models. We demonstrate the effectiveness of this combination function on real data sets and compare against several baselines including the state-of-the-art MLN learning algorithms.

The rest of the paper is organized as follows: first we introduce the necessary background and introduce the notations. Next we derive the gradients and present the algorithm

for learning RLR models. Finally, we conclude the paper by presenting our extensive experimental evaluations on standard SRL data sets and a NLP data set and outlining the areas for future research.

Background and Notation

In this section, we define our notation and provide necessary background for our readers to follow the rest of the paper. Throughout the paper, we assume *True* is represented by 1 and *False* is represented by 0.

Logistic Regression

Let Q be a Boolean random variable with range $\{1, 0\}$ whose value depends on a set $\{X_1, X_2, \dots, X_n\}$ of random variables. Logistic regression (McCullagh 1984) defines the conditional probability of Q given X_1, X_2, \dots, X_n as the sigmoid of a weighted sum of X_i s:

$$Prob(Q = 1 | X_1, \dots, X_n) = \sigma(w_0 + w_1 X_1 + \dots + w_n X_n) \quad (1)$$

where $\sigma(z) = 1/(1 + \exp(-z))$ is the sigmoid function.

Finite-Domain, Function-Free, First-Order Logic

A **population** is a finite set of objects. We assume for every object, there is a unique **constant** denoting that object. A **logical variable (logvar)** is typed with a population. A **term** is a logvar or a constant. We show logvars with lower-case letters (e.g., x), constants with upper-case letters (e.g., X), the population associated with a logvar x with Π_x , and the size/cardinality of the population with $|\Pi_x|$. A lower-case letter in bold represents a tuple of logvars (e.g., \mathbf{x}), and an upper-case letter in bold is a tuple of constants (e.g., \mathbf{X}).

An **atom** is of the form $Q(t_1, \dots, t_k)$ where Q is a functor and each t_i is a term. When $range(Q) = \{1, 0\}$, Q is a predicate. A **substitution** is of the form $\theta = \{\langle x_1, \dots, x_k \rangle / \langle t_1, \dots, t_k \rangle\}$ where x_i s are logvars and t_i s are terms. A **grounding** of an atom with logvars x_1, \dots, x_k is a substitution $\{\langle x_1, \dots, x_k \rangle / \langle X_1, \dots, X_k \rangle\}$ mapping each of its logvars to a constant in the population of the logvar. For a set \mathcal{A} of atoms, $\mathcal{G}(\mathcal{A})$ represents the set of all possible groundings for the atoms in \mathcal{A} . A **literal** is an atom or its negation. A **formula** φ is a literal, the conjunction of two formulae $\varphi_1 \wedge \varphi_2$, or a disjunction of two formulae $\varphi_1 \vee \varphi_2$. The application of a substitution $\theta = \{\langle x_1, \dots, x_k \rangle / \langle t_1, \dots, t_k \rangle\}$ on a formula φ is represented as $\varphi\theta$ and replaces each x_i in φ with t_i . An **instance** of a formula φ is obtained by replacing each logvar x in φ by one of the objects in Π_x . A **conjunctive formula** contains no disjunction. A **weighted formula (WF)** is a triple $\langle \varphi, w_T, w_F \rangle$ where φ is a formula and w_T and w_F are real numbers.

Relational Logistic Regression

Let $Q(\mathbf{x})$ be a Boolean atom whose probability depends on a set \mathcal{A} of atoms such that $Q \notin \mathcal{A}$. We refer to \mathcal{A} as the parents of Q . Let ψ be a set of WFs containing only atoms from \mathcal{A} , J be a function from groundings in $\mathcal{G}(\mathcal{A})$ to truth values, and $\theta = \{\mathbf{x}/\mathbf{X}\}$ be a substitution from logvars in \mathbf{x} to constants in \mathbf{X} . **Relational logistic regression (RLR)**

(Kazemi et al. 2014a) defines the probability of $Q(\mathbf{X})$ given J as follows:

$$Prob_\psi(Q(\mathbf{X}) = 1 | J) = \sigma \left(w_0 + \sum_{\langle \varphi, w_T, w_F \rangle \in \psi} w_T \cdot \eta_T(\varphi\theta, J) + w_F \cdot \eta_F(\varphi\theta, J) \right) \quad (2)$$

where w_0 is a bias/intercept, $\eta_T(\varphi\theta, J)$ is the number of instances of $\varphi\theta$ that are true with respect to J , and $\eta_F(\varphi\theta, J)$ is the number of instances of $\varphi\theta$ that are false with respect to J . Note that $\eta_T(\text{True}, J) = 1$. Also note that the bias can be considered as a WF whose formula is *True*. Following Kazemi et al. (Kazemi et al. 2014a), without loss of generality we assume the formulae of all WFs for RLR models are conjunctive.

Example 1 Let *Active*(p), *AdvisedBy*(s, p), and *PhD*(s) be three atoms representing respectively whether a professor is active, whether a student is advised by a professor, and whether a student is a PhD student. Suppose we want to define the conditional probability of *Active*(p) given the atoms $\mathcal{A} = \{\text{AdvisedBy}(s, p), \text{PhD}(s)\}$. Consider an RLR model with an intercept of -3.5 which uses only the following WF to define this conditional probability:

$$\psi = \{ \langle \text{AdvisedBy}(s, p) \wedge \text{PhD}(s), 1, 0 \rangle \}$$

According to this model, for an assignment J of truth values to $\mathcal{G}(\mathcal{A})$:

$$Prob_\psi(\text{Active}(P) = 1 | J) = \sigma(-3.5 + 1 \cdot \eta_T(\text{AdvisedBy}(s, P) \wedge \text{PhD}(s), J)),$$

where $\eta_T(\text{AdvisedBy}(s, P) \wedge \text{PhD}(s), J) = \#S \in \Pi_s$ s.t. $\text{AdvisedBy}(S, P) \wedge \text{PhD}(S)$ according to J , corresponding to the number of PhD students advised by P . When this count is greater than or equal to 4, the probability of P being an active professor is closer to one than zero; otherwise, the probability is closer to zero than one. Therefore, this RLR model represents “a professor is active if the professor advises at least 4 PhD students”.

With this background on Relational Logistic Regression, we introduce the Functional Gradient Boosting paradigm in the following section. This enables us to formulate a learning problem for RLR in which we learn both the structure and the parameters simultaneously.

Functional Gradient Boosting

We discuss **functional gradient boosting (FGB)** approach in the context of relational models. This approach is motivated by the intuition that finding many rough rules-of-thumb of how to change one’s probabilistic predictions locally can be much easier than finding a single, highly accurate model. Specifically, this approach turns the problem of learning relational models into a series of **relational function approximation** problems using the ensemble method of **gradient-based boosting**. This is achieved by the application of Friedman’s (Friedman 2001) gradient boosting to

SRL models. That is, we represent the conditional probability distribution as a weighted sum of regression models that are grown via a stage-wise optimization (Natarajan et al. 2012; Khot et al. 2011).

The conditional probability of an example y_i ¹ depends on its parents $\mathbf{x}_i = \text{parents}(y_i)$. The goal of learning is to fit a model $Prob(y | \mathbf{x}) \propto e^{\psi(y, \mathbf{x})}$, and can be expressed non-parametrically in terms of a potential function $\psi(y_i; \mathbf{x}_i)$:

$$Prob(y_i | \mathbf{x}_i) = \frac{e^{\psi(y_i; \mathbf{x}_i)}}{\sum_{y'} e^{\psi(y'; \mathbf{x}_i)}} \quad (3)$$

At a high-level, we are interested in successively approximating the function ψ as a sum of weak learners, which are relational regression clauses, in our setting. Functional gradient ascent starts with an initial potential ψ_0 and iteratively adds gradients Δ_i . After m iterations, the potential is given by $\psi_m = \psi_0 + \Delta_1 + \dots + \Delta_m$. Here, Δ_m is the **functional gradient** at episode m and is

$$\Delta_m = \eta_m \cdot E_{\mathbf{x}, y} \left[\frac{\partial}{\partial \psi_{m-1}} \log P(y | \mathbf{x}; \psi_{m-1}) \right], \quad (4)$$

where η_m is the learning rate. Dietterich *et al.* (Dietterich, Ashenfelter, and Bulatov 2004) suggested evaluating the gradient at every position in every training example and fitting a regression tree to these derived examples i.e., fit a regression tree h_m on the training examples $[(x_i, y_i), \Delta_m(y_i; x_i)]$. They point out that although the fitted function h_m is not exactly the same as the desired Δ_m , it will point in the same direction (assuming that there are enough training examples). Thus, ascent in the direction of h_m will approximate the true functional gradient.

Note that in the functional gradient presented in (4), the expectation $E_{\mathbf{x}, y}$ cannot be computed as the joint distribution $P(\mathbf{x}, y)$ is unknown. Instead of computing the functional gradients over the potential function, they are instead computed **pointwise** for each labeled training example i : $\langle \mathbf{x}_i, y_i \rangle$. Now, this set of local gradients become the training examples to learn a weak regression model that approximates the gradient Δ_m at stage m .

The functional gradient with respect to $\psi(y_i = 1; \mathbf{x}_i)$ of the likelihood for each example $\langle y_i, \mathbf{x}_i \rangle$ can be shown to be:

$$\frac{\partial \log Prob(y_i; \mathbf{x}_i)}{\partial \psi(y_i = 1; \mathbf{x}_i)} = \mathbb{I}(y_i = 1; \mathbf{x}_i) - P(y_i = 1; \mathbf{x}_i) \quad (5)$$

where \mathbb{I} is the indicator function, that is 1, if $y_i = 1$, and 0 otherwise. This expression is simply *the adjustment required to match the predicted probability with the true label of the example*. If the example is positive and the predicted probability is less than 1, this gradient is positive indicating that the predicted probability should move towards 1. Conversely, if the example is negative and the predicted probability is greater than 0, the gradient is negative driving the value the other way.

¹We use the term example to mean the grounded target literal. Hence $y_i = 1$ denotes that the grounding $Q(\mathbf{X}) = 1$ i.e., the grounded target predicate is true. Following standard Bayesian networks terminology, we denote the parents $\mathcal{A}(Q)$ to include the set of formulae ψ that influence the current predicate Q .

This elegant gradient expression might appear simple, but in fact, naturally and intuitively captures, example-wise, the general direction that the overall model should be grown in. The $I - P$ form of the functional gradients is a consequence of the sigmoid function as a modeling choice, and is a defining characteristic of FGB methods. As we show below, our proposed approach to RLR also has a similar form. The significant difference, however, is in the novel definition of the potential function ψ .

In prior work, **relational regression trees** (RRTs) (Blockeel 1999) were used to fit the gradient function Δ_m to the pointwise gradients for every training example. Each RRT can be viewed as defining several new feature combinations, one corresponding to each path from the root to a leaf. A key difference in our work is that we employ the use of weighted formulae (vector-weighted clauses², to be precise) as we explain later. From this perspective, our work is closer to the boosting MLN work that employed the use of weighted clauses. We generalize this by learning a weight vector per clause that allows for a more compact representation of the true and false instances of the formula. An example of a weighted clause is provided in Figure 1 where there are four clauses for predicting *advisedBy(A,B)*. Note that while we show the standard weighted clauses similar to a MLN, our weighted clauses have an important distinction - corresponding to each clause is a weight vector instead of a single scalar weight which captures the weight of true groundings, false groundings and the uninformed prior weights of the clause. The gradient-boosting that we develop in the next section builds upon these clauses and as mentioned earlier is similar to MLN boosting with the key difference being that instead of learning one weight per clause, we learn three weights in the vector.

The key intuition with boosting regression clauses is that each clause will define a new feature combination and the different clauses together capture the *latent relationships* that are learned from the data. While the final model itself is linear (as it is the sum of the weighted groundings of all the clauses), the clauses themselves define richer features thus allowing for learning a more complex model than a simple linear one. Figure 2 presents the schematic for boosting. The idea is that first a regression function (shown as a set of clauses) is learned from the training examples and these clauses are then used to determine the gradients (weights) of each example in the next iteration. The gradient is typically computed in the prior work as $I - P$. Once the examples are weighted, a new set of clauses are induced from them. These clauses are then considered together and the *regression values are added* when weighing the examples and the process is iterated.

There are several benefits of the boosting approach for learning RLR models. First, being a non-parametric approach (i.e., the model size is not chosen in advance), the number of parameters naturally grows as the number of training episodes increases. In turn, relational features as clauses are introduced only as necessary, so that a potentially large search space is not explicitly considered. Second, such

²We use formulae and clauses interchangeably from hereon.

$$\begin{aligned}
w_1 &: \text{professor}(B) \wedge \text{professor}(A) \rightarrow \text{advisedBy}(A, B) \\
w_2 &: \text{professor}(B) \wedge \neg \text{professor}(A) \wedge \text{Coauthor}(A, B) \rightarrow \text{advisedBy}(A, B) \\
w_3 &: \text{professor}(B) \wedge \neg \text{professor}(A) \wedge \neg \text{Coauthor}(A, B) \rightarrow \text{advisedBy}(A, B) \\
w_4 &: \neg \text{professor}(B) \rightarrow \text{advisedBy}(A, B)
\end{aligned}$$

Figure 1: Example of an RRC. The task was to predict if A is AdvisedBy B , given the relations of people at a university.

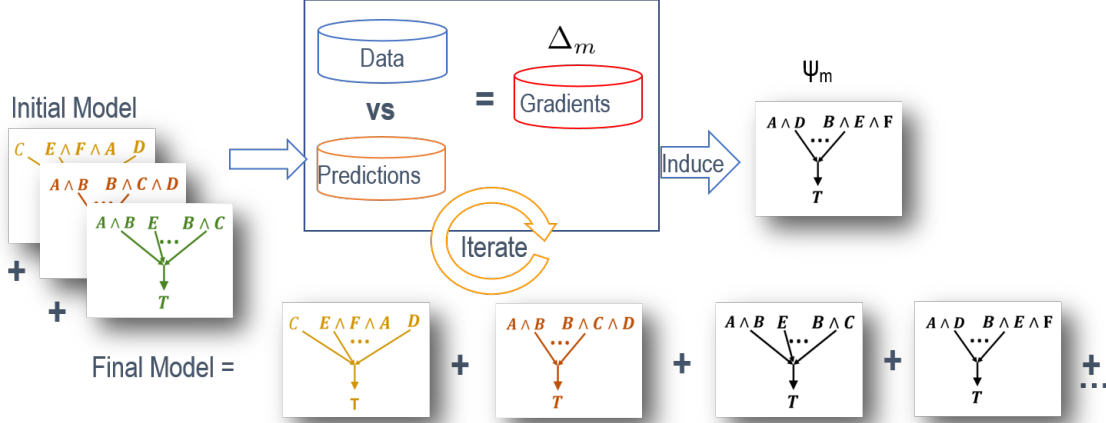


Figure 2: Relational Functional Gradient Boosting. This is similar to standard functional gradient boosting (FGB) where trees are induced stage-wise; the key difference is that these trees are relational regression clauses (RRCs).

an algorithm is fast and straightforward to implement. One could potentially employ any relational regression learner in the inner loop to learn several types of models. Third, as with previous relational models, the use of boosting for learning RLR models makes it possible to **learn the structure and parameters simultaneously** making them an attractive choice for learning from large scale data sets (Malec et al. 2016; Yang et al. 2017).

Functional Gradient Boosting for RLR

Preliminaries

Given the background on RLR and the gradient-boosting approach, we now focus on the learning task of RLR. Let us rewrite the conditional probability of an example y given weighted formulae $\langle \varphi_1, w_{T_1}, w_{F_1} \rangle, \dots, \langle \varphi_k, w_{T_k}, w_{F_k} \rangle$ corresponding parents J_1, \dots, J_k in the RLR model as:

$$\begin{aligned}
\text{Prob}(y = 1 \mid J_1, \dots, J_k) = & \\
& \sigma(w_0 + w_{T_1} \eta_T(\varphi_1 \theta, J_1) + w_{F_1} \cdot \eta_F(\varphi_1 \theta, J_1) \\
& + \dots \\
& + w_{T_k} \eta_T(\varphi_k \theta, J_k) + w_{F_k} \cdot \eta_F(\varphi_k \theta, J_k)) \quad (6)
\end{aligned}$$

where $\sigma(\cdot)$ is the sigmoid function. For example, let $y = \text{Popularity}(a)$ indicate the popularity of a professor a . Consider two formulae $\phi_1 = \text{Publication}(A, P)$ and $\phi_2 = \text{AdvisedBy}(A, S)$. The weights of the first formula control the influence of the number of publications on the popularity of the professor where $J_1 = \text{Publication}(a, P)$. Similarly the second formula controls the influence of

the number of students advised by the professor. For learning a model for RLR, we thus need to learn these clauses ϕ_i and their weights w_{T_i}, w_{F_i} (the parents are determined by the structure of the clause). Also, we can assume that the bias term w_0 can be part of the weight vectors for all the learned clauses. This allows a greedy approach that incrementally adds new clauses, such as FGB, to automatically update the bias term by learning w_0 for each new clause. Our learning problem can be defined as:

Given: A set of grounded facts (features) and the corresponding positive and negative grounded literals (examples)

To Do: Learn the set of formulae φ_i with their corresponding weight vector $\mathbf{w}_i = [w_0, w_1, w_2]$.

To simplify the learning problem, we introduce **vector-weighted clauses** (formulae), denoted as $\{\mathbf{w} : \text{Clause}\}$, that are a generalization of traditional weighted clauses with single weights. More specifically, our weighted clauses employ **three dimensions**, that is $\mathbf{w} = [w_0, w_1, w_2]^T$, where w_0 is a bias/intercept, w_1 is the weight over the satisfiable groundings of the current clause (analogous to w_{T_i}) and w_2 is the weight of the unsatisfiable groundings of the current clause (analogous to w_{F_i}). We also use a short hand notation t_i and f_i for the two grounding counts $\eta_T(\varphi_1 \theta, J_1)$ and $\eta_F(\varphi_1 \theta, J_1)$ respectively in Equation 6.

Example 2 Consider an RLR model for defining the conditional probability of $y = \text{Popularity}(a)$ which has only one

WF:

$$\langle \text{Publication}(A, B), w_T, w_F \rangle \equiv [w_0, w_1, w_2] : \text{Popularity}(A) := \text{Publication}(A, B)$$

Let $t_y = \sum_b \text{Publication}(a, b)$ be the number of instances of b for which $\text{Publication}(a, b)$ is true for the current grounding of y , and let $f_y = \sum_b (1 - \text{Publication}(a, b))$ be the number of instances of b for which $\text{Publication}(a, b)$ is false for the example y . Using vector-weighted clauses in the RLR model, we can compute

$$\text{Prob}(\text{Popularity}(a) \mid \text{Publication}(a, B)) = \sigma(w_0 + w_1 \cdot t_y + w_2 \cdot f_y), \forall a.$$

RFGB for RLR

Our goal is to learn the full structure of the model, which involves learning two concepts – the structure (formulae/clauses) and their associate parameters (the weight vectors). To adapt functional gradient boosting to the task of learning RLR, we map this probability definition over the parameter space w_0, w_1, w_2 to the functional space, ψ :

$$P(y_i = 1 \mid \mathbf{x}_i) = \sigma(\psi(y_i; \mathbf{x}_i)) = \sigma(w_0 + w_1 \cdot t_i + w_2 \cdot f_i) \quad (7)$$

Recall that in FGB, the gradients of the likelihood function with respect to the potential function are computed separately for each example. Correspondingly, the regression function ψ for the i -th example needs to be clearly defined and is:

$$\psi(y_i; \mathbf{x}_i) = w_0 + w_1 \cdot t_i + w_2 \cdot f_i. \quad (8)$$

The key difference to the existing gradient boosting methods for RDNs (Natarajan et al. 2012) and MLNs (Khot et al. 2011) is that the RLR learning algorithm needs to learn a weight vector per clause instead of a single weight.

Also, recall that while in traditional parametric gradient-descent, one would compute the parametric gradient over the loss function and iteratively update the parameters with these gradients, for gradient boosting, we first compute the functional gradients over the log-likelihood given by:

$$\Delta(y_i; \mathbf{x}_i) = \mathbb{I}(y_i = 1) - P(y_i = 1; \mathbf{x}_i) \quad (9)$$

where \mathbb{I} is the indicator function. As with other relational functional gradients (see Section Functional Gradient Boosting), this elegant expression naturally falls out when the log-likelihood of the sigmoid is differentiated with respect to the function. As before, the gradient is simply the adjustment required for the probabilities to match the observed value (y_i) in the training data for each example. Note that this is simply the outer gradient, that is, the gradient is computed for each example and a single vector-weighted clause needs to be learned for this set of gradients. While learning the clause itself, we must optimize a different loss function as we show next.

In order to generalize beyond the training examples, we fit a regression function ψ (which is essentially a *vector-weighted clause*) over the training examples such that the

squared error between $\psi(y_i; \mathbf{x}_i)$ and the functional gradient $\Delta(y_i; \mathbf{x}_i)$ is minimized over all the examples. The inner loop thus amounts to learning vector-weighted clauses such that we minimize the (regularized) squared error between the RLR model and the functional gradients over the n training examples:

$$\underset{w_0, w_1, w_2}{\text{minimize}} \sum_{i=1}^n (w_0 + w_1 t_i + w_2 f_i - \Delta(y_i; \mathbf{x}_i))^2 + \lambda (w_0^2 + w_1^2 + w_2^2), \quad (10)$$

where $\lambda > 0$ is a regularization parameter. In principle, λ can be chosen using a line search with a validation set when the size of the data sets are large. However, in our data sets, we only considered a few λ values from the set of $\{10^2, 10^{2.5}, 10^3, 10^{3.5}\}$ and chose to present the best λ corresponding to the test set.

Close inspection of the loss function above reveals that solving this optimization problem amounts to fitting **count features**: $\mathbf{c}_i = [1, t_i, f_i]$ for each grounded example i to the corresponding functional gradient, Δ_i . Note that the equation (10) can be viewed as a regularized least-squares regression problem to identify weights $\mathbf{w} = [w_0, w_1, w_2]^T$. The problem (10) can be written in vector form as

$$\min_{\mathbf{w}} \|\mathbf{C}\mathbf{w} - \Delta\|^2 + \lambda \|\mathbf{w}\|^2$$

where the i -th row of the count matrix C are the count features \mathbf{c}_i of the i -th example. This problem has a closed-form solution that can be computed efficiently:

$$\mathbf{w} = (\mathbf{C}^T \mathbf{C} + \lambda \mathbf{I})^{-1} \mathbf{C}^T \Delta. \quad (11)$$

The quantity $\mathbf{C}^T \mathbf{C}$ captures the **count covariance** across examples, while the quantity $\mathbf{C}^T \Delta$ captures the **count-weighted gradients**:

$$\mathbf{C}^T \mathbf{C} = \begin{bmatrix} n & \sum_{i=1}^n t_i & \sum_{i=1}^n f_i \\ \sum_{i=1}^n t_i & \sum_{i=1}^n t_i^2 & \sum_{i=1}^n t_i f_i \\ \sum_{i=1}^n f_i & \sum_{i=1}^n t_i f_i & \sum_{i=1}^n f_i^2 \end{bmatrix}, \quad (12)$$

$$\mathbf{C}^T \Delta = \begin{bmatrix} \sum_{i=1}^n \Delta_i \\ \sum_{i=1}^n t_i \Delta_i \\ \sum_{i=1}^n f_i \Delta_i \end{bmatrix}.$$

In this manner, functional gradient boosting enables a natural combination of conditionals over all the examples. This weight update forms the backbone of our approach: boosted relational logistic regression or B-RLR.

Algorithm for B-RLR

We outline the algorithm for boosted RLR (B-RLR) learning in Algorithm 1. We initialize the regression function with an uniform prior γ i.e. $F_0(y_i) = \gamma$ (line 2). Given the input training examples Y which correspond to the grounded instances of the target predicate y and the set of facts, i.e., the grounded set of all other predicates (denoted as X) in the domain, the goal is to learn the set of vector-weighted clauses that influence the target predicate.

Since there could potentially be multiple target predicates (when learning a joint model such as RDN where each influence relation is an RLR), we denote the current predicate as p . In the m^{th} iteration of functional-gradient boosting, we compute the functional gradients for these examples using the current model F_m and the parents of \mathbf{y} as per this model (line 7). Given these regression examples S_p , we learn a vector-weighted clause using FITREGRESSION. This function uses all the other facts X to learn the structure and parameters of the clause. We then add this regression function, $\hat{\psi}_m$ approximating the functional gradients to the current model, F_m . We repeat this over M iterations where M is typically set to 10 in our experiments.

Algorithm 1 Boosted Relational Logistic Regression (B-RLR) learning

```

1: function B-RLR( $Y, X, p$ )
2:    $F_0 := \gamma$ 
3:   for  $1 \leq m \leq M$  do ▷ M gradient steps
4:      $F_m := F_{m-1}$ 
5:     ▷ Compute gradients,  $\Delta_i$  for  $y_i \in Y_p$ 
6:      $S_p := \text{COMPUTEGRADIENTS}(Y_p, X, F_m)$ 
7:      $\hat{\psi}_m := \text{FITREGRESSION}(S_p, X, Y_p)$  ▷ Learn
vector-weighted regression clause
8:      $F_m := F_m + \hat{\psi}_m$  ▷ Update model
9:   end for
10:  return  $F_m$ 
11: end function

```

Next, we describe FITREGRESSION to learn vector-weighted clauses from input regression examples S , facts D and target predicate $p(x)$ in Algorithm 2. We initialize the vector-weighted clause with empty body and zero weights i.e. $[0, 0, 0]^T : y :- \emptyset$. We first create all possible literals that can be added to the clause given the current body (line 5). We use **modes** (Muggleton and De Raedt 1994) from inductive logic programming (ILP) to efficiently find the relevant literals here.

For each literal l in this set, we calculate the true and false groundings for the newly generated clause by adding the literal to the body (line 9). To perform this calculation, we ground the left hand side of the horn clause (i.e., the query literal) and count the number of groundings of the body corresponding to the query grounding. For instance if the grounding of the query is *advisedBy(John, Mary)* corresponding to *advisedBy(student, prof)*, then we count the number of instances of the body that correspond to *John* and *Mary*. If the body contains the publications in common, they are counted accordingly. If the body is about courses John took, they are counted correspondingly. This is similar to counting in any relational probabilistic model such as MLNs or BLPs. Following standard SRL models, we assume closed-world. This allows us to deduce the number of false groundings as the difference between the total number of possible groundings and the number of counted (positive) groundings.

We can then calculate the count matrix C_ℓ and weights \mathbf{w} as described earlier (line 11–12). We score each literal based

on the squared error and greedily pick the best scoring literal \hat{l} . We repeat this process till the clause reaches its maximum allowed length (set to 4 in our experiments).

Algorithm 2 Vector-weighted regression clause learning

```

1: function FITREGRESSION( $S, D, y$ )
2:   where  $S = \{(y_i, \Delta_i)\}$ 
3:   body :=  $\emptyset$ ;  $\mathbf{w} := [0, 0, 0]^T$  ▷ Initialize empty
clause
4:   while  $\text{len}(\text{body}) \leq C$  do
5:      $L := \text{POSSIBLELITERALS}(p(x), \text{body})$  ▷
Generate potential literals
6:     for  $\ell \in L$  do ▷ Score each literal
7:       clause := ‘ $y :- \text{body} \wedge \ell$ .’
8:       for  $x_i \in X$  do ▷ Calculate groundings per
example
9:          $t_i, f_i = \text{CALCULATEGROUNDINGS}(y_i,$ 
 $D, \text{clause})$ 
10:      end for
11:       $C_\ell := \text{CREATECOUNTMATRIX}(\{t_i, f_i\})$ 
12:       $\mathbf{w}(\ell) := (C^T C + \lambda I)^{-1} C^T \Delta$ .
13:       $\text{score}(\ell) := \text{SCOREFIT}(\mathbf{w}(\ell), \Delta)$ 
14:    end for
15:     $\hat{\ell} := \arg \min_{\ell} \text{score}(\ell)$ 
16:     $\mathbf{w} := \mathbf{w}(\hat{\ell})$ 
17:    body := body  $\wedge \hat{\ell}$ 
18:  end while
19:  return  $\mathbf{w} : y :- \text{body}$ .
20: end function

```

To summarize, given a target, the algorithm computes the gradient for all the examples based on the expression $I - P$. Given these gradients, the inner loop searches over the possible clauses such that the MSE is minimized. The resulting vector-weighted clauses are then added to the set of formula and are then used for the subsequent steps of gradient computations. The procedure is repeated until convergence or a preset number of formulae are learned. The search for the most optimal clause can be guided by experts by providing relevant search information as modes (Muggleton and De Raedt 1994). The overall procedure is similar to RDNs and MLNs with two significant differences - the need for multiple weights in the clauses and correspondingly the different optimization function inside the inner loop of the algorithm.

Given that we have outlined the B-RLR algorithm in detail, we now turn our focus to empirical evaluation of this algorithm.

Experiments and Results

Our experiments will aim to answer the following questions in order to demonstrate the benefits of B-RLR:

- Q1** How does functional gradient boosting perform when compared to traditional learning approaches for clauses and weights?
- Q2** How does the boosted method perform compared to a significant feature engineered logistic regression approach?

Domains	Sample Features
WorkedUnder (IMDB) 5 features constructed	count_genres_acted, count_movies_acted
AdvisedBy (UWCS) 8 features constructed	count_publications, count_taughtby
Female (Movie lens) 8 features constructed	count_movies, average_ratings
Cancer (SmCaFr) 3 features constructed	no_of_friends, no_of_friends_smoke
Faculty (WebKB) 4 features constructed	count_project, count_courseta

Table 1: This table shows the number of rules used by AGG-RLR for each data set as well as the features with high weights as picked by LR

- Q3** How does boosting RLR compare to other relational methods?
- Q4** How sensitive is the behavior of the proposed approach with respect to the regularization constant, λ ?

Methods Considered

We now compare our B-RLR approach to: (1) the AGG-LR approach, which is standard logistic regression (LR) using the relational information, (2) the ILP-RLR approach where rules are learned using a logic learner, followed by weight learning for the formulae, and (3) MLN-B, which is a state-of-the-art boosted MLN structure learning method. We evaluate our approach on 1 synthetic data set and 4 real world data sets. Table 1 shows the sample aggregate (Relational) features constructed with the highest weights as generated by AGG-LR.

A natural question to ask is the comparison of our method against the recently successful Boosted Relational Dependency Networks (Natarajan et al. 2012) (bRDN) method. We do not consider this comparison for two important reasons - first is that the MLN-B has already been compared against bRDN in the original work and the conclusion was that they were nearly on par in performance in all the domains while bRDN is more efficient due to the use of existentials instead of counts when grounding clauses. Consequently, the second reason is that since our AGG-LR approach heavily employs counts, we considered the best learning method that employs counts as an aggregator, namely the MLN-B method. Our goal is not to demonstrate that AGG-LR is more effective than the well-known MLN-B or the bRDN approaches, but to demonstrate that boosting RLR does not sacrifice performance of learners and that RLR can be boosted as effectively as other relational probabilistic models.

In contrast to our approach, which performs parameter and structure learning *simultaneously*, the ILP-RLR baseline performs these steps *sequentially*. More specifically, we use PROGOL (Muggleton 1995; Muggleton 1997) for structure learning, followed by relational logistic regression for parameter learning. Table 3 shows the number of rules that were learned for each data set by PROGOL. Table 3 also shows some sample rules with the highest coverage scores as generated by PROGOL.

To keep comparisons as fair as possible, we used the following protocol: while employing MLN-B, we set the maximum number of clauses to 3, the beam-width to 10 and maximum clause length to 4. Similar configurations were adopted in our clause-learning setting. Gradient steps for MLN-B and B-RLR were picked as per the performance.

Data Sets

Smokes-Cancer-Friends: This is a synthetic data set, where the goal is to predict who has cancer based on the friends network of individuals and their observed smoking habits. The data set has three predicates: Friends, Smokes and Cancer. We evaluated the method over the Cancer predicate using the other predicates with 4-fold cross-validation and $\lambda = 10^{3.5}$.

UW-CSE: The UW-CSE data set (Richardson and Domingos 2006) was created from the University of Washington’s Computer Science and Engineering department’s student database and consists of details about professors, students and courses from 5 different subareas of computer science (AI, programming languages, theory, system and graphics). The data set includes predicates such as Professor, Student, Publication, AdvisedBy, HasPosition, TaughtBy etc., Our task is to learn, using the other predicates, to predict the AdvisedBy relation between a student and a professor. There are 4,106,841 possible AdvisedBy relations out of which only 3380 are true. We employ 5-fold cross validation where we learn from four areas and predict on the other area with $\lambda = 10^{3.5}$ in our reported results.

IMDB: The IMDB data set was first used by Mihalkova and Mooney (Mihalkova and Mooney 2007) and contains five predicates: Actor, Director, Movie, Genre, Gender and WorkedUnder. We predict the WorkedUnder relation between an actor and director using the other predicates. Following (Kok and Domingos 2009), we omitted the four equality predicates. We set $\lambda = 10^3$ and employed 5-fold cross-validation using the folds generation strategy suggested by Mihalkova and Mooney in (Mihalkova and Mooney 2007) and averaged the results.

WebKB: The WebKB data set was first created by Craven et al. (Craven et al. 1998) and contains information about department webpages and the links between them. It also contains the categories for each web-page and the words within each page. This data set was converted by Mihalkova and Mooney (Mihalkova and Mooney 2007) to contain only the category of each web-page and links between these pages. They created the following predicates: Student, Faculty, CourseTA, CourseProf, Project and SamePerson from these web-pages. We evaluated the method over the Faculty predicate using the other predicates and we performed 4-fold cross-validation where each fold corresponds to one university with set λ set as 10^2 .

Movie Lens: This is the well-known Movielens data set (Harper and Konstan 2015) containing information of about 940 users, 1682 movies, the movies rated by

Data Sets	Target	Types	Predicates	neg:pos Ratio
Sm-Ca-Fr	Cancer	1	3	1.32
IMDB	WorkedUnder	3	6	13.426
UW-CSE	AdvisedBy	9	12	539.629
WebKB	Faculty	3	6	4.159
Movie Lens	FemaleGender	7	6	2.702

Table 2: Details of relational domains used in our experiments. These data sets have high ratios of negative to positive examples, which is a key characteristic of relational data sets.

each user containing 79,778 user-movie pairs, and the actual rating the user has given to a movie. It contains predicates: Age, Genre, Occupation, Year, Ratings and Gender. In our experiments, we ignored the actual ratings and only considered whether a movie was rated by a user or not. Also, since Gender can take only two values, we convert the $Gender(person, gender)$ predicate to a single argument predicate $FemaleGender(person)$. We learned B-RLR models for predicting $FemaleGender$ using 5-fold cross-validation with $\lambda = 10^3$.

A key property of these relational data sets is the large number of negative examples. This is depicted in Table 2, which shows the size of various data sets used in our experiments. This is because, in relational settings, a vast majority of relations between objects are not true, and the number of negative examples far outnumbers the number of positive examples. In these data sets, simply measuring accuracy or log-likelihood can be misleading. Hence, we use metrics which are reliable in imbalanced setting like ours.

Results

We present the results of our experiments in Figures 3 and 4, which compare the various methods on two metrics: area under the ROC curve (AUC-ROC) and area under the Precision-Recall curve (AUC-PR) respectively. From these figures, certain observations can be made clearly.

First, the proposed B-RLR method is on par or better than most methods across all data sets. On deeper inspection, it can be seen that the state-of-the-art boosting method for MLNs appears to be more mixed at first glance in ROC-space while B-RLR is generally better in PR-space. In addition, in the WebKB, MovieLens and Smokes-Cancer-Friends domain where we learn about a unary predicate, the performance is significantly better. This yields an interesting insight: RLR models can be **natural aggregators** over the associated features. As we are in the unary predicate setting (which corresponds to predicting an attribute of an object), the counts of the instances of the body of the clause simply means aggregating over the values of the body. This is typically done in several different ways such as mean, weighted mean or noisy-or (Natarajan et al. 2008). We suggest the use of logistic function with counts as an alternative aggregator that seems effective in this domain and we hypothesize its use for many relational tasks where aggregation can yield to natural models. In contrast, MLNs only employ counts as their features, while RLR allows for a more complex ag-

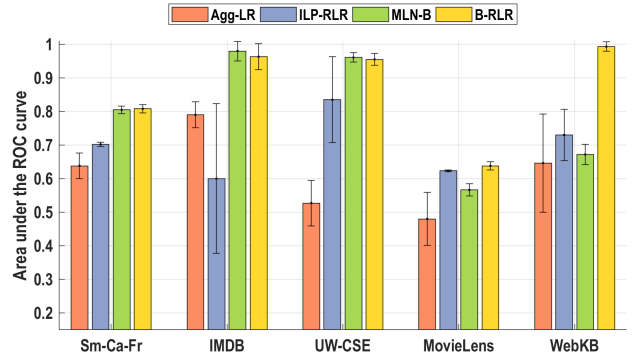


Figure 3: Comparing the area under the ROC curve for the proposed B-RLR approach to (1) standard logistic regression with relational information (AGG-LR), (2) an approach where rules are learned using a logic learner followed by weight learning (ILP-RLR), and (3) state-of-the-art MLN with boosted structure learning (MLN-B).




Figure 4: Comparing the area under the Precision-Recall (PR) curve for the proposed B-RLR approach to (1) standard logistic regression with relational information (AGG-LR), (2) an approach where rules are learned using a logic learner followed by weight learning (ILP-RLR), and (3) state-of-the-art MLN with boosted structure learning (MLN-B).

gregation within the sigmoid function that can use count features in its inner loop. Validating this positive aspect of RLR models remains an interesting future research direction. These results help in answering **Q3** affirmatively: that B-RLR is on par or significantly better than MLN-B in all domains.

Next, from the figures, it can be observed that B-RLR significantly outperforms AGG-LR in several domains. At the outset, this may not be surprising since relational models have been shown to outperform non-relational models. However, the features that are created for the AGG-LR model are the count features of the type defined in the original RLR work and are more expressive than the standard features of propositional models. This result is particularly insightful as the B-RLR model that uses count features, predicates and their combinations themselves in a formula is far more expressive than simple aggregate features. This allows us to an-

Target (Data Set)	Sample rules generated for ILP-RLR using PROGOL
WorkedUnder (IMDB) 6 rules generated	WorkedUnder(A, B) \Leftarrow isa(B, director), isa(A, actor), movie(C, A), movie(C, B). WorkedUnder(A, B) \Leftarrow genre(B, C), gender(A, male).
AdvisedBy (UWCS) 16 rules generated	AdvisedBy(A, B) \Leftarrow hasPosition(B, E), inPhase(A, D), publication(C, A), publication(C, B). AdvisedBy(A, B) \Leftarrow hasPosition(B, D), inPhase(A, E), publication(F, A).
Female (Movie Lens) 7 rules generated	Female(A) \Leftarrow tmpRatingArg1(B, A), tmpRatingArg2(B, C), genre(C, g4). Female(A) \Leftarrow age(A, 4), occupation(A, o14).
Cancer (SmCaFr) 3 rules generated	Cancer(a) \Leftarrow friends(b, a), friends(b, c), smokes(c). Cancer(a) \Leftarrow smokes(a).
Faculty (WebKB) 6 rules generated	Faculty(A) \Leftarrow courseProf(B, A), courseTA(B, C). Faculty(A) \Leftarrow courseProf(B, A), project(C, A), samePerson(A, A).

Table 3: This table shows the number of rules used by ILP-RLR for each data set as well as the rules with the highest $\|P - N\|$ coverage as returned by PROGOL.

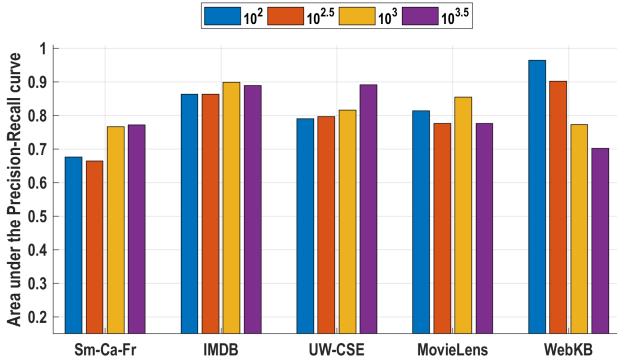


Figure 5: Sensitivity of the proposed B-RLR approach is analyzed by comparing the Area under the Precision-Recall (PR) curve as λ changes.

swer **Q2** strongly and affirmatively, in that the proposed approach is significantly better than an engineered (relational) logistic regression approach.

Finally, comparing the proposed approach to a two-step approach of learning clauses followed by the corresponding weights; B-RLR appears to be significantly better in both PR-space as well as ROC-space than ILP-RLR. Furthermore, B-RLR also has the distinct advantage of simultaneous parameter and structure learning, thus avoiding a costly structure search compared to the ILP-based approach. Hence, **Q1** can be answered: that B-RLR model outperforms ILP-based two-stage learning in all the regions. We see from Figure 5 that **Q4** is answered affirmatively: As long as λ is within the set $\{10^2, 10^{2.5}, 10^3, 10^{3.5}\}$, our algorithm is not sensitive in most domains. In addition, our parameter λ has a nice intuitive interpretation: they reflect and incorporate the high class imbalance that exists for real-world domains where this is an important practical consideration.

We used paired t-test with p-values=0.05 for determining the statistical significance. From the Figures 3 & 4, across most domains, we observe that, B-RLR has tighter error bounds compared to the baselines in majority of domains indicating lower variance and subsequently higher generalization performance.

Table 4 reports the training time taken in seconds by each method averaged over all the folds in every domain. Tim-

Target (Data set)	Learning Time (seconds)			
	AGG-LR	ILP-RLR	MLN-B	B-RLR
WorkedUnder (IMDB)	67.39	158.02	6.99	10.95
AdvisedBy (UWCS)	110.58	65.058	18.40	24.71
Female (Movie Lens)	91.27	78.57	6.66	16.51
Cancer (SmCaFr)	35.18	3.26	120.73	90.47
Faculty (WebKB)	51.73	2.18	5.68	5.86

Table 4: Comparing learning time (in seconds) for the proposed B-RLR approach to (1) standard logistic regression with relational information (AGG-LR), (2) an approach where rules are learned using a logic learner followed by weight learning (ILP-RLR), and (3) state-of-the-art MLN with boosted structure learning (MLN-B). Learning time includes time to learn the structure, counting the satisfied (or unsatisfied) groundings and weight learning.

ings reported For AGG-LR include time taken for propositional feature construction and weight learning using WEKA tool. For ILP-RLR it includes the total time taken to learn rules, count satisfied instances, and learn weights for the rules accordingly. The two boosted approaches timings are reported from the full runs. The results show that the methods are comparable across all the domains - in the domains where boosted methods are faster than the other baselines, grounding of the entire data set caused the increased time for the baselines. Conversely, in the other domains, repeated counting of boosting increased the time in two of the five domains. The results indicate that the proposed B-RLR approach does not sacrifice efficiency (time) for effectiveness (performance).

In summary, our proposed boosted approach appears to be promising across all a diversity of relational domains, with the potential for scaling up relational logistic regression models to large data sets.

Conclusions

We considered the problem of learning relational logistic regression (RLR) models using the machinery of functional-gradient boosting. To this end, we introduce an alternative interpretation of RLR models that allows us to consider both the true and false groundings of a formula within a single equation. This allowed us to learn vector-weighted clauses

that are more compact and expressive compared to standard boosted SRL models. We derived gradients for the different weights, and outlined a learning algorithm that learned first-order features as clauses and the corresponding weights simultaneously. We evaluated the algorithm on standard data sets, and demonstrated the efficacy of the learning algorithm.

There are several possible extensions for future work – currently, our method learns a model for a single target predicate deterministically. As mentioned earlier, it is possible to learn a joint model across multiple predicates in a manner akin to learning a relational dependency network (RDN). This can yield a new interpretation for RDNs based on combining rules. Second, learning from truly hybrid data remains a challenge for SRL models in general, and RFGB in particular. Finally, given the recent surge of sequential decision-making research, RLR can be seen as an effective function approximator for relational Markov decision processes (MDPs); employing this novel B-RLR model in the context of relational reinforcement learning can be an exciting and interesting future research direction.

References

- [Blokkeel 1999] Blokkeel, H. 1999. Top-down induction of first order logical decision trees. *AIC* 12(1-2).
- [Craven et al. 1998] Craven, M.; DiPasquo, D.; Freitag, D.; McCallum, A.; Mitchell, T.; Nigam, K.; and Slattery, S. 1998. Learning to extract symbolic knowledge from the world wide web. In *AAAI*, 509–516.
- [Dietterich, Ashenfelter, and Bulatov 2004] Dietterich, T.; Ashenfelter, A.; and Bulatov, Y. 2004. Training CRFs via gradient tree boosting. In *ICML*.
- [Fatemi, Kazemi, and Poole 2016] Fatemi, B.; Kazemi, S. M.; and Poole, D. 2016. A learning algorithm for relational logistic regression: Preliminary results. *arXiv preprint arXiv:1606.08531*.
- [Friedman 2001] Friedman, J. 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* 29.
- [Getoor and Taskar 2007] Getoor, L., and Taskar, B. 2007. *Introduction to Statistical Relational Learning*. MIT Press.
- [Harper and Konstan 2015] Harper, F., and Konstan, J. 2015. The MovieLens datasets: History and context. *ACM Trans. Interact. Intell. Syst.* 5(4):19:1–19:19.
- [Heckerman, Meek, and Koller 2007] Heckerman, D.; Meek, C.; and Koller, D. 2007. Probabilistic entity-relationship models, PRMs, and plate models. *Introduction to statistical relational learning* 201–238.
- [Huynh and Mooney 2008] Huynh, T., and Mooney, R. 2008. Discriminative structure and parameter learning for Markov logic networks. In *ICML*, 416–423. ACM.
- [Kazemi and Poole 2018] Kazemi, S. M., and Poole, D. 2018. ReINN: A deep neural model for relational learning. In *AAAI*.
- [Kazemi et al. 2014a] Kazemi, S.; Buchman, D.; Kersting, K.; Natarajan, S.; and Poole, D. 2014a. Relational logistic regression. In *KR*.
- [Kazemi et al. 2014b] Kazemi, S.; Buchman, D.; Kersting, K.; Natarajan, S.; and Poole, D. 2014b. Relational logistic regression: The directed analog of markov logic networks. In *AAAI Workshop*.
- [Kersting and De Raedt 2007] Kersting, K., and De Raedt, L. 2007. Bayesian logic programming: Theory and tool. In Getoor, L., and Taskar, B., eds., *An Introduction to Statistical Relational Learning*.
- [Khot et al. 2011] Khot, T.; Natarajan, S.; Kersting, K.; and Shavlik, J. 2011. Learning Markov logic networks via functional gradient boosting. In *ICDM*, 320–329. IEEE.
- [Kimmig et al. 2012] Kimmig, A.; Bach, S.; Broecheler, M.; Huang, B.; and Getoor, L. 2012. A short introduction to probabilistic soft logic. In *NIPS Workshop*, 1–4.
- [Kok and Domingos 2009] Kok, S., and Domingos, P. 2009. Learning Markov logic network structure via hypergraph lifting. In *ICML*, 505–512. ACM.
- [Koller 1999] Koller, D. 1999. Probabilistic relational models. In *ILP*, 3–13. Springer.
- [Malec et al. 2016] Malec, M.; Khot, T.; Nagy, J.; Blasch, E.; and Natarajan, S. 2016. Inductive logic programming meets relational databases: An application to statistical relational learning. In *ILP*.
- [McCullagh 1984] McCullagh, P. 1984. Generalized linear models. *EJOR* 16(3):285–292.
- [Mihalkova and Mooney 2007] Mihalkova, L., and Mooney, R. 2007. Bottom-up learning of Markov logic network structure. In *ICML*, 625–632. ACM.
- [Muggleton and De Raedt 1994] Muggleton, S., and De Raedt, L. 1994. Inductive logic programming: Theory and methods. *Journal of Logic Programming* 19:629–679.
- [Muggleton 1995] Muggleton, S. 1995. Inverse entailment and prolog. *New generation computing* 13(3-4):245–286.
- [Muggleton 1997] Muggleton, S. 1997. Learning from positive data. In *ILP*, 358–376. Berlin, Heidelberg: Springer Berlin Heidelberg.
- [Natarajan et al. 2008] Natarajan, S.; Tadepalli, P.; Dietterich, T.; and Fern, A. 2008. Learning first-order probabilistic models with combining rules. *Annals of Mathematics and Artificial Intelligence* 54(1-3):223–256.
- [Natarajan et al. 2012] Natarajan, S.; Khot, T.; Kersting, K.; Gutmann, B.; and Shavlik, J. 2012. Gradient-based boosting for statistical relational learning: The relational dependency network case. *Machine Learning* 86(1):25–56.
- [Neville and Jensen 2007] Neville, J., and Jensen, D. 2007. Relational dependency networks. In Getoor, L., and Taskar, B., eds., *Introduction to Statistical Relational Learning*. MIT Press. 653–692.
- [Poole et al. 2014] Poole, D.; Buchman, D.; Kazemi, S.; Kersting, K.; and Natarajan, S. 2014. Population size extrapolation in relational probabilistic modelling. In *SUM*, 292–305. Springer.
- [Raedt et al. 2016] Raedt, L.; Kersting, K.; Natarajan, S.; and Poole, D. 2016. Statistical relational artificial intelligence:

Logic, probability, and computation. *Synthesis Lectures on AI and ML* 10(2):1–189.

[Richardson and Domingos 2006] Richardson, M., and Domingos, P. 2006. Markov logic networks. *ML* 62(1):107–136.

[Taskar et al. 2007] Taskar, B.; Abbeel, P.; Wong, M.; and Koller, D. 2007. Relational Markov networks. *Introduction to statistical relational learning* 175–200.

[Yang et al. 2017] Yang, S.; Korayem, M.; AlJadda, K.; Grainger, T.; and Natarajan, S. 2017. Combining content-based and collaborative filtering for job recommendation system: A cost-sensitive statistical relational learning approach. *KBS* 136:37–45.