

Combining Gradient Boosting Machines with Collective Inference to Predict Continuous Values

Iman Alodah

Computer Science Department
Purdue University
West Lafayette, Indiana 47906
Email: ialodah@purdue.edu

Jennifer Neville

Computer Science and Statistics Departments
Purdue University
West Lafayette, Indiana 47906
Email: neville@cs.purdue.edu

Abstract

Gradient boosting of regression trees is a competitive procedure for learning predictive models of continuous data that fits the data with an additive non-parametric model. The classic version of gradient boosting assumes that the data is independent and identically distributed. However, relational data with interdependent, linked instances is now common and the dependencies in such data can be exploited to improve predictive performance. Collective inference is one approach to exploit relational correlation patterns and significantly reduce classification error. However, much of the work on collective learning and inference has focused on discrete prediction tasks rather than continuous. In this work, we investigate how to combine these two paradigms together to improve regression in relational domains. Specifically, we propose a boosting algorithm for learning a collective inference model that predicts a continuous target variable. In the algorithm, we learn a basic relational model, collectively infer the target values, and then iteratively learn relational models to predict the *residuals*. We evaluate our proposed algorithm on a real network dataset and show that it outperforms alternative boosting methods. However, our investigation also revealed that the relational features interact together to produce better predictions.

Introduction

Supervised learning methods use observed data to fit predictive models. The generalization error of these prediction models can be decomposed into two main sources: *bias* and *variance*. In general, for a single model, there is a tradeoff between bias and variance. However, ensemble methods attempt to resolve this issue by learning a combination of models. *Boosting* is an ensemble approach that creates a strong predictor from a combination of multiple weak predictors. The method aims to reduce errors due to bias *and* variance to produce better predictions. In this work, we consider the gradient boosting algorithm (Friedman 2000), which fits an additive model to independent and identically distributed (i.i.d.) data in a forward stage-wise manner.

However, many current datasets of interest do not consist of i.i.d. instances, but instead contain structured, relational, network information. For example, in Facebook

users are connected to friends, and information about the friends (e.g., demographics, interests) can be helpful to boost prediction results. As such, there has been a great deal of recent work focused on the development of *collective inference* methods that can exploit relational correlation among class label values of interrelated instances to improve predictions (Chakrabarti, Dom, and Indyk 1998; Neville and Jensen 2000; Taskar, Abbeel, and Koller 2002; Lu and Getoor. 2003a; Macskassy and Provost 2007; Sen et al. 2008; McDowell, Gupta, and Aha 2009). Collective inference methods typically learn local relational models and then apply them to *jointly* classify unlabeled instances in a network. Since network data often exhibit high correlation among related instances, this results in improved classification performance compared to traditional i.i.d. methods.

In this paper, we investigate the combination of gradient boosting (GB) and collective inference (CI) to predict continuous-valued target variables in partially-labeled networks. While boosting has previously been considered in the context of statistical relational learning (Natarajan et al. 2010; Khot et al. 2011; Natarajan et al. 2012; Hadiji et al. 2015; Khot et al. 2015), our work is different in two key aspects. First, we consider the problem of regression rather than discrete classification. Second, our proposed method uses *residuals* at every stage of the gradient boosting algorithm, and uses collective inference to exploit correlations in the the residuals during prediction. Specifically, we propose a *modified gradient boosting* (MGB) algorithm for learning a collective inference model for relational regression. In the MGB algorithm, we learn a basic relational model, collectively infer the target values, and then iteratively learn relational models to predict the *residuals*. We evaluate MGB on a real network dataset and show that it outperforms alternative GB methods. We also conduct some ablation studies that show the importance of the various relational features. The results indicate that relational features involving the neighbors' residuals were not as helpful as features that use the (current) boosted predictions.

The paper is organized as follows. We first review the basic GB and CI methods used to develop our algorithm. Then, we describe the proposed MGB algorithm in detail, including both learning and inference. Next, we describe the experimental setup and show the results. Finally, we discuss related work and conclude.

Background

Gradient Boosting Machine (GB):

Gradient boosting machines (Friedman 2000) are a powerful machine-learning technique that combines a set of weak learners into a single strong learner by fitting an additive model in a forward stage-wise manner. Given a random target variable Y and a set of random attribute variables $X = X_1, X_2, \dots, X_p$, the goal is to use the training samples $\langle y_i, \mathbf{x}_i \rangle$ to estimate a function $F(\mathbf{x})$ that maps \mathbf{x} to y . An estimate of the true underlying function $F^*(\mathbf{x})$ is learned by minimizing the expected value of a specified loss function L :

$$\arg \min_F E_{y, \mathbf{x}}[L(y, F(\mathbf{x}))] \quad (1)$$

The approximation function is of the form:

$$F_M(\mathbf{X}; \rho_m, \theta) = \sum_{m=0}^M \rho_m h_m(\mathbf{X}; \theta) \quad (2)$$

where $h_m(\mathbf{X}; \theta)$ is a function of the input variables \mathbf{x} (parametrized by θ) and ρ_m is a multiplier. $\{h_m(\mathbf{X})\}_1^M$ are incremental functions (or *boosts*) and h_0 is the initial guess. The goal is to minimize the objective function using a greedy stage-wise approach:

$$(\rho_m, \theta_m) = \arg \min_{\rho, \theta} \sum_{i=1}^N L(y_i, [F_{m-1}(\mathbf{x}_i) + \rho \cdot h(\mathbf{x}_i; \theta)]) \quad (3)$$

where L is the loss function. To solve equation 3 we choose $h(\mathbf{x}; \theta)$ that is highly correlated with the negative gradient by solving:

$$\theta_m = \arg \min_{\theta} \sum_{i=1}^N [\hat{y}_i - h(\mathbf{x}_i; \theta)]^2 \quad (4)$$

where \hat{y}_i is the residual:

$$\hat{y}_i = y_i - F(\mathbf{x}_i) = - \left[\frac{\partial L(y_i, F(\mathbf{x}_i))}{\partial F(\mathbf{x}_i)} \right]_{F(\mathbf{X})=F_{m-1}(\mathbf{X})} \quad (5)$$

The multiplier ρ_m is estimated by applying line search:

$$\rho_m = \arg \min_{\rho} \sum_{i=1}^N L(y_i, [F_{m-1}(\mathbf{x}_i) + \rho \cdot h(\mathbf{x}_i; \theta_m)]) \quad (6)$$

At each stage the model is updated as follows:

$$F_m(\mathbf{X}) = F_{m-1}(\mathbf{X}) + \rho_m h(\mathbf{X}; \theta_m) \quad (7)$$

One choice for the models h are regression trees. In this work, we fit the negative gradient with a regression tree. Algorithm 1 shows pseudocode for general gradient boosting. Note that lines 4-5 refer to learning a regression tree model, if regression trees are used as the basic model h . ρ is computed for each region in the regression tree as suggested by (Friedman 2000) and it is actually done by taking the median of the region.

Input : data samples $\{\langle y_i, \mathbf{x}_i \rangle\}_1^N$, iterations M
Output: Model F_M which has the form of Eq. 2

```

1 Compute initial guess  $F_0(X)$ , e.g., for least square
  regression use  $F_0 = \bar{y}$ .
2 for  $m=1$  to  $M$  do
3   Compute the gradient for each example with Eq. 5.
   For regression  $\hat{y}_i = y_i - F_{m-1}(\mathbf{x}_i)$ .
4   Estimate parameters  $\theta_m$  with Eq. 4.
5   Compute the step length  $\rho_m$  with Eq. 6.
6   Update the model  $F_m$  with Eq. 7.
7 end
8 return  $F_M$ 

```

Algorithm 1: Gradient Boosting

Iterative Collective Inference:

Collective classification is jointly used to infer the unknown class labels in a partially-labeled attributed graph $G = (V, E)$. Each node $v_i \in V$ has an associated class label and attributes $\langle y_i, \mathbf{x}_i \rangle$. In *transductive* collective classification, the goal is to learn a model from the subset of nodes in G that are labeled (which we refer to as Y_k), and then apply collective inference to make predictions for the remaining (unlabeled) nodes in the graph $Y_u = Y - Y_k$. To exploit the relational information, we use a *relational* learner model that is trained using the objects' features and relational features R_f of their neighbors.

Algorithm 2 shows the iterative classification algorithm ICA we use for collective inference. Iterative classification algorithm is one of several algorithms that performs collective classification. In this work, we will use the same process to perform iterative collective *regression*. In other words, our main focus in this work is to use ICA to predict values \tilde{Y} in continuous space which uses \mathcal{N}_i the neighbors of node v_i in the prediction process. The relational features that we compute are aggregations of continuous class values (e.g., average and median).

Collective Inference Gradient Boosting For Continuous Variables

Problem Statement:

Many relational and social network datasets exhibit correlation among the class labels values of linked nodes. While there is a large body of work focused on learning collective classification models for predicting discrete target variables, there is relatively less work focused on continuous class labels. In this paper we focus on developing a boosting algorithm for predicting a continuous target variable Y .

Definition 1 *Transductive collective classification in relation data with continuous class labels:*

Given a partially labeled graph $G = (V, E, \mathbf{X}, Y)$ where V is a set of nodes, E is a set of edges, \mathbf{X} is set of features, and $y_i \in Y$ is a continuous target value for each node $v_i \in V$. The nodes V are split into a *known* set V_k with labels Y_k and an *unknown* set V_u s.t., $V = V_k + V_u$. The goal is to learn

Input : Graph $G = (V = V_k + V_u, E, \mathbf{X}, Y_k), F$
 V contains two types of nodes: V_k have known class labels and V_u are unknown; \mathbf{x}_i are the features of node v_i ; ; F is the learned model.

- 1 **for** each v_i in V_u **do**
- 2 Compute relational features (Rf_i) using v_i 's observed neighbors (i.e., $y_j \in Y_k$ s.t. $v_j \in \mathcal{N}_i$)
- 3 Construct feature vectors $A_i = [x_i, Rf_i]$
- 4 $\tilde{y}_i \leftarrow F(A_i)$
- 5 **end**
- 6 **repeat** *Iterative inference*
- 7 Generate random ordering O over V_u
- 8 **for** each v_i in O **do**
- 9 Compute relational features (Rf_i) using known labels and current predictions for \mathcal{N}_i (i.e., y_j for $v_j \in V_k$ and \tilde{y}_j for $v_j \in V_u$)
- 10 Construct feature vectors $A_i = [x_i, Rf_i]$
- 11 $\tilde{y}_i \leftarrow F(A_i)$
- 12 **end**
- 13 **until** *converges or number of iterations have elapsed*

Algorithm 2: Iterative Classification Algorithm ICA

a collective classification model from the induced subgraph over V_k and apply it to G to predict class label values for Y_u .

Model:

In this work, we develop a gradient boosting algorithm that uses collective inference (CI) in the learning phase. Since we focus on regression, we modify the gradient boosting algorithm that minimizes squared loss as described in Friedman (2000) and combine it with CI. Our Modified Gradient Boosting (MGB) algorithm is described in Alg. 3. The algorithm fits a stronger non-parametric function as in Eq. 2—the learning is guided by CI over the continuous target variable Y , which uses the full model F_m and the residuals \hat{y}_i that are predicted using the weak learned models $\{h\}_1^m$. In the application phase, we use CI to collectively infer the values of Y_u using the induced relational gradient-boosted model. Algorithm 4 shows how we apply the ICA algorithm in our setting. Note we use \hat{y}_i refer to the residual of the node v_i and \tilde{y}_i refer to the estimation of the class label of the node v_i (i.e., $\hat{y}_i = y_i - \tilde{y}_i$).

The base model of our MGB algorithm is a regression tree. Our model is an additive model that learns a tree at each stage, and then the final result is the sum of the predictions from each tree. We made one additional modification to the original squared loss gradient boosting regression in Friedman (2000). Since our data exhibits a skewed class distribution, we use the median instead of the average to compute predictions in the terminal nodes of each tree. Using the median prevents the predictions from being skewed towards the extreme values in the distribution.

Algorithm 3 starts by fitting a regression tree with a specified depth l . Then the input graph is divided into two separate sets for CI: 20% are set as unknown in Y_{tu} and 80% are

Input : Graph $G_{tr} = (V_k, E_k, \mathbf{X}_k, Y_k)$, the induced subgraph over nodes with known labels;
 l : max number of nodes in regression tree;
 M : number of boosting iterations;
 t : number of \tilde{Y} estimations.

Output: Model F_M which has the form of Eq. 2.

- 1 Compute relational features Rf using Y_k
- 2 Construct feature vectors $\mathbf{A}_0 = [\mathbf{X}_k, Rf] \forall v \in V_k$
- 3 Fit regression tree $F_0(\mathbf{A}_0) = h_0(\mathbf{A}_0, l)$ to predict Y_k
- 4 **for** t iterations **do**
- 5 Randomly split G_{tr} into 80% to use for Y_{tk} and 20% for Y_{tu} .
- 6 $\tilde{Y}_{tu} = \text{ICA2}(V_{tk}, V_{tu}, E_k, \mathbf{X}_k, Y_{tk}, F_0)$
- 7 **end**
- 8 $\tilde{y}_i =$ average of \tilde{y}_{ti} for $t \in 1..t$
- 9 **for** $m = 1$ to M **do**
- 10 Compute residuals $\hat{y}_i = y_i - \tilde{y}_i$ for all $v \in V_k$
- 11 Compute relational features Rf using $Y_k, \hat{Y}_k, \tilde{Y}_k$
- 12 Construct feature vectors $\mathbf{A}_m = [\mathbf{X}_k, Rf]$
- 13 Fit regression tree $h_m(\mathbf{A}_m, l)$ to predict \hat{Y}
- 14 Update the final model:
 $F_m(\mathbf{A}_{[0..m]}) = F_{m-1}(\mathbf{A}_{[0..m-1]}) + h_m(\mathbf{A}_m)$
- 15 **for** t iterations **do**
- 16 Randomly split G_{tr} into 80/20% for Y_{tk} / Y_{tu}
- 17 $\tilde{Y}_{tu} = \text{ICA2}(V_{tk}, V_{tu}, E_k, \mathbf{X}_k, Y_{tk}, F_m)$
- 18 **end**
- 19 $\tilde{y}_i =$ average of \tilde{y}_{ti} for $t \in 1..t$
- 20 **end**
- 21 **return** F_M

Algorithm 3: Modified Gradient Boosting (MGB)

set as known in Y_{tk} . The division of labeled/unlabeled can be tuned experimentally. These two sets are then used in the CI process when ICA2 is applied. We repeat the CI multiple times to ensure that every node in the input graph is set as unknown at least once. This can be achieved by choosing a random order of G_{tr} at every iteration then dividing it into the two groups as discussed above. The multiple predictions for each node are averaged afterwards (line 8).

The algorithm then enters the loop that fits a set of models that constitute the *residual* models. In line 10, the algorithm computes the residuals (i.e., the negative gradient). In line 11, the relational features are updated. Note that each model in $\{F_m\}_0^m$ has a potentially different set of relational features. The features are specific to each model F_m because they depend on the target value of that model (i.e., class label or residual). The next section describes the set of relational features in more detail.

Once the relational features are recalculated, the algorithm learns a new tree model over the residuals then it applies ICA2 with the new complete model $F_{m-1} + h$, where h is the regression tree learned on the current iteration. This approach aims to exploit the correlation in the residuals for more accurate predictions. As before, we need to repeat the inference process multiple times to ensure that each node gets at least one prediction from CI. The predictions are then averaged to be used in the next iteration of residual calcula-

```

Input : Graph  $G = (V = V_k + V_u, E, \mathbf{X}, Y_k)$ ,  $F_M$ ,  $itr$ 
           $F_M$  is the learned model;
           $itr$ : number CI iterations.
Output:  $\tilde{y}$  estimations of unknown labels  $Y_u$ 

1 # compute residuals of known set
  for  $m \in [1, M]$  do
2   for each  $v_i$  in  $V_k$  do
3     Compute relational features ( $Rf_{im}$ ) using  $v_i$ 's
       observed neighbors (i.e.,
4      $y_j \in Y_k$  s.t.  $v_j \in \mathcal{N}_i$ )
       Construct feature vector  $A_{im} = [x_i, Rf_{im}]$ 
5      $\tilde{y}_{im} \leftarrow F_{m-1}(A_{im})$ 
6     Compute true residual  $\hat{y}_{im} = y_i - \tilde{y}_{im}$ 
7   end
8 end
9 # initialize predicted class/residuals for unknown
   $\mathcal{Y}_i \leftarrow [] \forall v_i \in V_u$ 
10 for  $m \in [0, M]$  do
11   for each  $v_i$  in  $V_u$  do
12     Compute relational features ( $Rf_{im}$ ) using  $v_i$ 's
       observed neighbors' class and residuals (i.e.,
13      $[y_j, \{\hat{y}_{im}\}_1^M]$  s.t.  $y_j \in Y_k \wedge v_j \in \mathcal{N}_i$ )
       Construct feature vector  $A_{im} = [x_i, Rf_{im}]$ 
14     if  $m = 0$  then  $\tilde{y}_i \leftarrow h_m(A_{im}); \mathcal{Y}_i[m] = \tilde{y}_i$ 
15     else  $\hat{y}_{im} \leftarrow h_m(A_{im}); \mathcal{Y}_i[m] = \hat{y}_{im}$ 
16   end
17 end
18 # collective inference over unknown set
  for  $t \in [0, itr]$  do
19   Generate random ordering  $O$  over  $V_u$ 
20   for each  $v_i$  in  $O$  do
21     for  $m \in [0, M]$  do
22       Compute relational features ( $Rf_{im}$ ) using
         known labels/residuals and current
         predictions (i.e.,  $[y_j, \{\hat{y}_{im}\}_1^M]$  for  $v_j \in V_k$ 
         and  $\mathcal{Y}_j$  for  $v_j \in V_u$ )
23       Construct feature vector  $A_{im} = [x_i, Rf_{im}]$ 
24       if  $m = 0$  then  $\tilde{y}_i \leftarrow h_m(A_{im}); \mathcal{Y}_i[m] = \tilde{y}_i$ 
25       else  $\hat{y}_{im} \leftarrow h_m(A_{im}); \mathcal{Y}_i[m] = \hat{y}_{im}$ 
26     end
27   end
28 end
29 return  $\sum_0^m \mathcal{Y}_{im} \forall v_i \in V_u$ 

```

Algorithm 4: ICA2 for MGB

tion. The number of specified iterations/models can also be tuned experimentally using a validation set to prevent overfitting.

ICA in algorithm 4 predicts not only the class label but also the residuals predicted by each sub model. The first loop from line 1 to 8 computes the true residuals for the known nodes. This information is need in the next loop from line 10 to 17 where we compute initial prediction for all the unknown nodes. Finally, in line 18 to 28 we apply collective inference to predict the class label and the residuals. Out final class label prediction is the sum of all the class label prediction and the residuals predicted in line 18 to 28.

Features:

Along with non-relational features, we use four different relational features. The first relational feature Rf_1 records the median of the neighbors' target values. The initial value of this feature is the median of the target values of the known neighbors and during the collective inference it is calculated using the known values and the estimated values over the neighbors:

$$Rf_{1j} = \text{median}\{\mathcal{N}_j(Y)\} \quad (8)$$

where Y are known or estimated values for the neighbors (\mathcal{N}_j) of node v_j .

The features $\{Rf_{2m}\}_0^M$ record the median of the neighbors' target values at each boosting stage m . The initial model F_0 predicts the class, so the feature will be the median of the target values. The models $m = 1 \dots M$ predict the residuals, so the associated feature values will be the median of the residuals computed at that stage:

$$Rf_{2mj} = \text{median}\{\mathcal{N}_j(\hat{Y}_m)\} \quad (9)$$

Here \hat{Y}_m refers to the known or estimated residuals at stage/model m . Note that at stage 0 we can view the full target value as the residual if we consider the initial estimation as 0.

The features $\{Rf_{3m}\}_0^M$ record the median of the residuals of the boosting stage $m - 1$, but it is used to train the model at stage m .

$$Rf_{3mj} = \text{median}\{\mathcal{N}_j(\hat{Y}_{m-1})\} \quad (10)$$

Here \hat{Y}_{m-1} are the known or estimated residuals at the stage/model $m - 1$.

Finally, the feature Rf_{4j} records the instance v_j 's target value estimated by the complete model so far (i.e., F_{m-1}). This feature is used in stage m to learn a new model:

$$Rf_{4j} = F_{m-1}(A_j) \quad (11)$$

Experimental Evaluation

Hypothesis:

Our hypothesis is that, when using gradient boosting to learn a relational regression model, exploiting relational correlation—of both predicted target values and residuals—can outperform learning using fixed relational features. In other words, we want to answer the following question:

Q1: Does using collective inference in the learning phase of the gradient boosting improves its prediction of continuous values compared to both gradient boosting without any relational features (GB) and gradient boosting with relational features but no collective inference (RGB)?

Datasets:

We evaluated our proposed algorithm on the IMDB data. We considered a network of movies that are linked if they share a producer, and the target label is the movie's total revenue. The largest connected component of this dataset has about 10770 nodes. Each movie has 29 genre features and we added the four relational features discussed above. Some statistics of the dataset are shown in Table 1.

Table 1: IMDB Dataset Characteristics

IMDB	
Characteristic	Value
Number of nodes	10770
Number of edges	212460
Labels correlation coefficient	0.22699
Average degree	19.727
Average clustering coefficient	0.53947
Number of non-relational features	29
Number of relational features	4
Target label	Continuous value from 30 to 1054×10^6

Baseline Models:

We compare our proposed MGB model to two different baselines. The first natural choice is the gradient boosting machine (GB) as proposed in Friedman (2000), which does not exploit the relational nature of network data. We use the 29 non-relational features to learn the model. The second baseline is a simple relational version of the gradient boosting machine (RGB) which is described next.

Relational Gradient Boosting: The RGB model is trained as regression tree with gradient boosting—using additional relational features. Along with the 29 non-relational features that we use to learn GB, we add one more relational feature that records the median of the known neighbors’ target labels Y_k . This is a simple relational model in which the attributes of an object depend on the class label of that object as well as the target labels of objects one link away (Jensen, Neville, and Gallagher 2004). The relational feature is computed as described in Eq. 8.

Experimental setup:

Training/Test Splits Generation We divide the network into 5 folds, where each is 20% of the data. Then we set the training set to be $\{20\%, 40\%, 60\%, 80\%\}$ of the data and the testing set to $(100\% - \text{training set})$.

Test Procedure We use 5-fold cross-validation for evaluation. The learning phase of the algorithms is implemented as described in Algorithm 3. For prediction, we apply the learned model using the ICA collective inference algorithm for RGB and the ICA2 collective inference algorithm for MGB. The accuracy of the predictions are evaluated with the root mean square error (RMSE).

The algorithms have different parameters that need to be specified in advance. For the tree-based gradient algorithm, we have two main parameters: (1) The number of boosting iterations M —which we set to 10, and (2) the max size of the regression tree l —which we set to 5. Both parameters can be tuned better by using a validation set and avoid over fitting the training set, but we leave this to future work.

We have another parameter that is related to collective inference. From the literature of iterative collective classification, 50 iteration (itr) are sufficient. Since we have continuous values instead of discrete, we assessed the convergence

of the estimations over the iterations and found that the estimations converge at a rate similar to discrete classification. Thus, we set $itr = 50$.

Our MGB algorithm has two additional parameters. The number of times t that inference is repeated to get an estimate for all the nodes. We set $t = 3$. Moreover we randomly selected 20% (80%) of the training set for the unknown (known) subset for each trial t .

Results and Discussion:

Our experiments show that incorporating collective inference in the learning phase improves performance—since the continuous target labels are correlated and the residuals (error) are also correlated. Figure 1 shows how our approach (MGB) outperforms the the baseline algorithms. Our algorithm achieves lower RMSE values when tested using different training set sizes (i.e., data proportion). We test the algorithms first by training using a random 20% of the network and apply the induced model on the remaining 80%. We then increase the proportions of the network allocated to training through to 80%. The performance of the MGB algorithm is consistently better than the other methods. The results also show that a larger portion of the network is labeled, the performance increases for all models (lower RMSE value) and the added gain provided by collective inference decreases.

One of the main factors that affect the performance is autocorrelation between the labels, the autocorrelation between the residuals and finally the correlation between the residuals and the true or estimated target values. In IMDB dataset, the labels are autocorrelated and the residuals are also autocorrelated. Feature Rf_1 and feature Rf_4 are the most effective features of the four suggested features. Feature Rf_1 utilizes the autocorrelation in the target label and using it alone gives good results. Using feature Rf_1 alone is comparable to how we would interpret the work of (Natarajan et al. 2012) for regression, note that both their task and the type of the target variable are different than ours. As for feature Rf_4 , it exploits the correlation between the residuals and the estimated target value. Using this feature alone gives good results as well, however, combining them gives the best performance. Figure 2 shows the performance of the MGB with only the relational feature Rf_4 and without it. The figure gives good indication that both Rf_4 and Rf_1 are needed for better performance. As for the relational feature Rf_1 , Figure 3 shows performance comparison when we implement MGB with only Rf_1 , all feature(MGB) and MGB without Rf_4 . It is also important to note that both features Rf_2 and Rf_3 are computed using the residuals. The performance of MGB does not get better when using only these features and might get even worse when rely on these features alone for prediction.

The level of relational dependence impacts performance, including (1) correlation between the labels of linked nodes, (2) correlation between the residuals of linked nodes, and (3) correlation between the residuals and the true or estimated target values. In the IMDB data, the labels and the residuals are correlated across links. Features Rf_1 and Rf_4 are the most effective features of the four suggested features. Feature Rf_1 utilizes the relation correlation in the target label and using



Figure 1: Performance of GB, RGB, and MGB on IMDB.

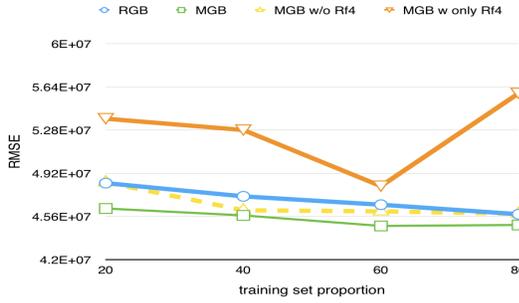


Figure 2: Performance of MGB learned using only Rf_4 and MGB learned with all the features but Rf_4 (on IMDB).

it alone gives good results. Using feature Rf_1 alone is comparable to how the Natarajan et al. (Natarajan et al. 2012) method would be implemented for regression. Note that in their paper, both their tasks and the type of the target variable are different from ours.

As for feature Rf_4 , it exploits the correlation between the residuals and the estimated target value. Using this feature alone produces good results as well, however, combining them results in the best performance. Figure 2 shows MGB performance with only the relational feature Rf_4 and without it. The results indicates that both Rf_4 and Rf_1 are needed to achieve best performance.

As for the relational feature Rf_1 , Figure 3 compares performance when we implement MGB with only the Rf_1 feature, all features (MGB), and MGB without Rf_4 . It is important to note that both features Rf_2 and Rf_3 are computed using the residuals. The performance of MGB does not improve when using only these features and might even degrade when rely on these features alone for prediction. Although the error is correlated with the residuals in our dataset, it is difficult to estimate the residuals using collective inference.

The accuracy of the algorithm is sensitive to the manner in which the features are initialized for learning in each model. We set the initial relational features for the models to be equal to the median of the known neighbors’ target values for Rf_1 and Rf_4 and zero for Rf_2 and Rf_3 . If a given node does not have any known neighbors, all the features are set to zero.

Related Work

Boosting for statistical relational learning (SRL) has been studied in previous work such as (Natarajan et al. 2012) and

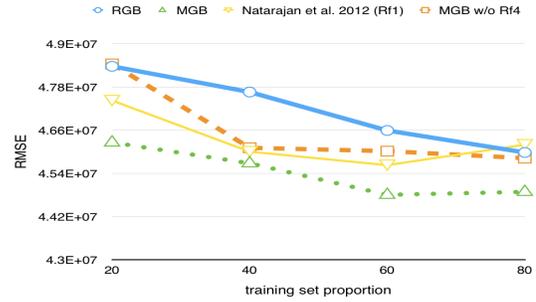


Figure 3: Performance of MGB compared to implementation of Natarajan et al. (2012) for regression (on IMDB).

(Natarajan et al. 2010). Natarajan et al. (2012) uses gradient boosting to learn relational and non-parametric functions that approximate a joint probability distribution over multiple random variables in *relational dependency networks* (RDNs). The algorithm learns the structure and parameters of the RDN models simultaneously. Since RDNs can be used for collective classification, Natarajan et al. include other query predicates in the training data while learning the model for the current query and then apply the learned mode collectively. The authors compare the boosted RDNs (RDN-B) against Markov logic networks (Richardson and Domingos 2006) and basic RDNs on two kinds of tasks: entity resolution and information extraction. The performance of RDN-B was significantly better compared to RDN and MLN for most datasets. However, the authors don’t report a comparison between using RDN-B with and without collective inference.

Hadji et al. (2015) introduce non-parametric boosted Poisson Dependency Networks (PDNs) using gradient boosting. Among other objectives in their work, the authors performed collective prediction, however, they again did not compare performance with and without collective inference.

In Khot et al. (2015; 2011), the authors used gradient-based boosting for MLNs. The authors derived functional gradients for the pseudo-likelihood and the boosted MLNs outperformed the non-boosted version. However, the boosted models were not used for collective inference.

We note that all the related work above focused on classification tasks with discrete labels rather than regression.

Conclusion

In this work we investigated the use of boosting to learn a collective regression model. Our results show that gradient boosting, combined with collective inference, results in improves performance compared to gradient boosting in a relational model that does not use collective inference. This indicates that using the *residuals* in the relational model, and including CI in the boosting process, can produce better predictions.

The work reported in this paper is different from other research on using boosting in SRL methods in several ways. First, our MGB method performs regression instead of classification. Second, we compared the performance of gradient boosting using simultaneous statistical estimates about the same variables for a set of related data instances (i.e.,

in collective inference) with the classic version of gradient boosting that considers a fixed set of features. We also compared to a simple relational model that does not use CI in the learning phase. Finally, we show that using collective inference to estimate the residuals is difficult and that the most effective features are the ones that does not use neighbors' residuals, but instead use their current boosted predictions.

Acknowledgments

This research is supported by NSF contract number IIS-1149789 and NSF Science & Technology Center grant CCF-0939370. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

References

- Chakrabarti, S.; Dom, B.; and Indyk, P. 1998. Enhanced hypertext categorization using hyperlinks. *SIGMOD Rec.* 27(2):307–318.
- Friedman, J. H. 2000. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*.
- Hadiji, F.; Molina, A.; Natarajan, S.; and Kersting, K. 2015. Poisson dependency networks: Gradient boosted models for multivariate count data. *Mach. Learn.* 100(2-3):477–507.
- Jensen, D.; Neville, J.; and Gallagher, B. 2004. Why collective inference improves relational classification. In *In Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 593–598.
- Khot, T.; Natarajan, S.; Kersting, K.; and Shavlik, J. 2011. Learning markov logic networks via functional gradient boosting. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining, ICDM '11*, 320–329. Washington, DC, USA: IEEE Computer Society.
- Khot, T.; Natarajan, S.; Kersting, K.; and Shavlik, J. 2015. Gradient-based boosting for statistical relational learning: The markov logic network and missing data cases. *Mach. Learn.* 100(1):75–100.
- Lu, Q., and Getoor, L. 2003a. Link-based classification. In *In Proceedings of the 20th International Conference on Machine Learning (ICML)*, 496–503.
- Macskassy, S. A., and Provost, F. 2007. Classification in networked data: A toolkit and a univariate case study. *J. Mach. Learn. Res.* 8:935–983.
- McDowell, L. K.; Gupta, K. M.; and Aha, D. W. 2009. Cautious collective classification. *J. Mach. Learn. Res.* 10:2777–2836.
- Natarajan, S.; Khot, T.; Kersting, K.; Gutmann, B.; and Shavlik, J. 2010. Boosting relational dependency networks. In *In Proc. of the Int. Conf. on Inductive Logic Programming (ILP)*.
- Natarajan, S.; Khot, T.; Kersting, K.; Gutmann, B.; and Shavlik, J. 2012. Gradient-based boosting for statistical relational learning: The relational dependency network case. *Mach. Learn.* 86(1):25–56.
- Neville, J., and Jensen, D. 2000. Iterative classification in relational data. In *In Proceedings of the Workshop on Statistical Relational Learning, 17th National Conference on Artificial Intelligence*, 42–49.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1-2):107–136.
- Sen, P.; Namata, G. M.; Bilgic, M.; Getoor, L.; Gallagher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI Magazine* 29(3):93–106.
- Taskar, B.; Abbeel, P.; and Koller, D. 2002. Discriminative probabilistic models for relational data. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence, UAI'02*, 485–492. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.