

Graph Based Classification Methods Using Inaccurate External Classifier Information

S Sundararajan*

S Sathiya Keerthi†

October 9, 2009

Abstract

In this paper we consider the problem of collectively classifying entities where relational information is available across the entities. In practice inaccurate class distribution for each entity is often available from another (external) classifier. For example this distribution could come from a classifier built using content features or a simple dictionary. Given the relational and inaccurate external classifier information, we consider two graph based settings in which the problem of collective classification can be solved. In the first setting the class distribution is used to fix labels to a subset of nodes and the labels for the remaining nodes are obtained like in a transductive setting. In the other setting the class distributions of *all* nodes are used to define the *fitting* function part of a graph regularized objective function. We define a generalized objective function that handles both the settings. Methods like harmonic Gaussian field and local-global consistency (LGC) reported in the literature can be seen as special cases. We extend the LGC and weighted vote relational neighbor classification (WvRN) methods to support usage of external classifier information. We also propose an efficient least squares regularization (LSR) based method and relate it to information regularization methods. All the methods are evaluated on several benchmark and real world datasets. Considering together speed, robustness and accuracy, experimental results indicate that the LSR and WvRN-extension methods perform better than other methods.

1 Introduction

Traditionally classifiers are built using only *local* features of individual entities such as web pages or images. Relational classifiers also make use of relational information that exist across the entities. For example, local features of a web page could be collection of keywords that appear in the title or page content and useful relational information could be scores computed from presence/absence of inlinks and/or outlinks, similarity of page structure, url etc. In relational classification prob-

lems, a collection of entities is represented as a graph or a network of nodes. Each node represents an entity with its set of local features, class attribute information and edge weight encodes any relational information that exists between those connected nodes. In its general setup the graph contains zero or more labeled nodes and one or more unlabeled nodes and the goal is to label a set of unlabeled nodes (test nodes). This problem is either solved by treating it as an induction problem where a model (possibly statistical) model is learnt using training data (labeled nodes) and used to classify future data, or, as a transduction problem where classification is needed(done) only on the test nodes. The key aspect of this relational classification problem is making *collective* inference of the class labels, that is, labels of all the nodes are obtained simultaneously. See [5], [4] and references there for more details.

In this paper we consider a related relational learning problem where, instead of a subset of labeled nodes, we have inaccurate external label/class distribution information for each node. This problem arises in many web applications. Consider, for example, the problem of identifying pages about Public works, Court, Health, Community development, Library etc. within the web site of a particular city. The link and directory relations contain useful signals for solving such a classification problem. Note that this relational structure will be different for different city web sites. If we are only interested in a small number of cities then we can afford to label a number of pages in each site and then apply transductive learning using the labeled nodes. But, if we want to do the classification on hundreds of thousands of city sites, labeling on all sites is expensive and we need to take a different approach. One possibility is to use a selected set of content dictionary features together with the labeling of a small random sample of pages from a number of sites to learn an inaccurate probabilistic classifier, e.g., logistic regression. Now, for any one city web site, the output of this initial classifier can be used to generate class distributions for the pages in the site, which can then be used together with the

*Yahoo! Labs, Bangalore, India, ssrajan@yahoo-inc.com

†Yahoo! Research, Santa Clara, CA, selvarak@yahoo-inc.com

relational information in the site to get accurate classification.

The *problem of doing relational learning together with externally available class distribution information* can be solved by modifying existing transductive methods. The problem has been discussed within broad sets of techniques such as *denoising, relaxation labeling, metric labeling* (see [3] and references there) etc., as well as within recent specific techniques such as the Gaussian Field Harmonic Function (GFHF) method [9] and the Information Regularization (IR)[2], dual IR (DIR) [6] methods. We also consider two more methods, viz., the Local-Global Consistency (LGC) method [8] and the probabilistic Weighted vote Relational Neighbor (WvRN) Classification method with Relaxation Labeling [4]. The main aim of this paper is to take a few techniques (the last five specific methods mentioned above) that are popularly discussed and used in recent relational learning literature, extend them as needed and compare in different settings to see which ones are most effective in solving our problem. The proposed extensions and method include supporting external classifier information for the LGC and WvRN methods, and using least squares (LS) divergence measure (referred to as the LSR method) as opposed to the KL-divergence measure used in the IR methods. We also establish the relation between the LSR method and the IR methods.

Our problem may be solved in two different settings. In the first setting we select a subset of nodes for which we have high confidence in the class labels and fix the labels of these nodes. For example we may select the nodes based on thresholding its probability or decision function score above a certain value. We can treat the remaining nodes as unlabeled and solve it as a transduction problem in which relational information is used to *propagate* labels from the labeled set to the unlabeled set through the connections according to their strengths. In the second setting we make use of the external class distribution of all the nodes fully and no nodes are treated as unlabeled. For different methods this is done in different ways.

Detailed experiments on four benchmark datasets indicate that the second solution setting (full use of class distribution information) is better. Further experiments (in setting 2) on real world shopping domain datasets clearly demonstrate that significant performance gain can be achieved with the proposed methods. Considering speed, robustness and accuracy together, the experimental results indicate that the LSR and WvRN extension methods are better than the other methods.

The paper is organized as follows. Section 2 formulates the problem and describes the two solution settings. The methods are grouped as function estimation

and probability distribution based methods; details on how they are modified for the two settings are given in sections 3 and 4. In section 5 we give experimental results, and summarize key observations in section 6.

2 Solution Settings.

In this section we present the statement of the problem and discuss two solution settings in which the problem can be solved. Then we briefly mention the methods and their extensions considered in this paper.

2.1 Notations and Problem Formulation. Assume that we have a graph G with vertices V and edges E . Let the vertices (nodes) v_i , $i \in N$ where $N = \{1, \dots, n\}$ represent the entities. Let the edges $e_{i,j}$ (where $i, j \in N$) encode relational information between node v_i and v_j , with weight $w_{i,j}$. Let \mathbf{W} denote the weight matrix with $w_{i,j}$ as its elements. Let Δ be the graph Laplacian matrix defined as $\Delta = \mathbf{D} - \mathbf{W}$ [9] and let its unnormalized and normalized versions be denoted as: $\mathbf{L}_{un} = \Delta$ and $\mathbf{L}_{norm} = \mathbf{D}^{-\frac{1}{2}} \Delta \mathbf{D}^{-\frac{1}{2}} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}}$ respectively. \mathbf{D} is a diagonal matrix with i th diagonal element defined as $d_{ii} = \sum_j w_{i,j}$ and d_{ii} measures degree of i th node. In a traditional problem formulation, labeling of v_i corresponds to specifying its class label c_i where $c_i \in \{1, \dots, K\}$ and K is the number of classes. A related quantity is the vector $\mathbf{y}_i = \delta_{k,c_i}$ where δ is the Kronecker delta function. Let $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]$. Alternatively, one can specify the class distribution information \mathbf{p}_i where $\mathbf{p}_i = [p_{i,1}, \dots, p_{i,K}]^T$ with $\sum_k p_{i,k} = 1$. Let $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_n]$. The case where only class label information (c_i , \mathbf{y}_i and \mathbf{Y}) are given may be viewed as a special case of the class distribution view with $\mathbf{P} = \mathbf{Y}$. Therefore without loss of generality we assume that we are given \mathbf{P} and we can derive the class label for node v_i as $c_i = \underset{k}{\operatorname{argmax}} p_{i,k}$. Thus, given \mathbf{P} , we can obtain derived labels c_i and obtain the corresponding performance (e.g., accuracy) $Perf(\mathbf{P})$. In this paper we are mainly concerned with effective use of an initial \mathbf{P} (that is available from some external means) and using it together with the relational information to do better. In such a case we call $Perf(\mathbf{P})$ as the *initial performance* of the classifier. The classification problem of using relational information (\mathbf{W}) and initial class distribution (\mathbf{P}) can be loosely stated as follows: given $(G, \mathbf{W}, \mathbf{P})$, find $\hat{\mathbf{P}}$ such that $Perf(\hat{\mathbf{P}})$ is better than $Perf(\mathbf{P})$. Given the above formulation we consider two settings in which the problem can be solved.

2.2 Solution Setting 1. In the first setting the external class distribution information available with all the nodes (\mathbf{P}) is first used to select a subset of la-

beled nodes, \mathbf{S} . We treat the remaining nodes as unlabeled. Then, for each $i \in \mathbf{S}$ we fix the label of node i as $\arg \max_k p_{i,k}$. Once this is done, the problem can be solved using any standard graph based transductive or semi-supervised learning method. However, while conventional semi-supervised learning or transduction problem settings usually assume that \mathbf{p}_i or \mathbf{y}_i is accurate, this is not the case here. The selection of subset of nodes is an important aspect and will be addressed later when we discuss specific methods.

2.3 Solution Setting 2. In the second setting the external class distribution information of *all* the nodes is used in the solution process. Exactly how this is done will be detailed in later sections.

2.4 Existing Methods and Extensions. Graph based transduction methods fall in one of two groups: (1) those which are based on function estimation and (2) those which are based on probability distribution estimation. In this paper we consider the following methods from each group: the Gaussian Field Harmonic Function (GFHF) method [9] and the Local-Global Consistency (LGC) method [8] from the first group and the Information Regularization (IR) [2], dual IR (DIR) [6] methods and the probabilistic Weighted vote Relational Neighbor (WvRN) Classification method with Relaxation Labeling [4] from the second group. While all these five methods can easily solve the problem in the first setting, the second setting is (briefly) discussed in the literature only for GFHF, IR and DIR methods. In the following sections we extend LGC and WvRN methods to handle the second setting. We also propose an efficient least squares regularization (LSR) based method as compared to the KL-divergence based IR methods and relate these methods.

3 Function Estimation Methods

Graph based classification methods in this group find the classification function by minimizing the function:

$$(3.1) \quad Q(\mathbf{F}) = Q_{GR}(\mathbf{F}) + CQ_{datafit}(\mathbf{F}, \mathbf{Y})$$

where C is a positive regularization constant. The objective function consists of two terms. The first term known as the *graph regularization* (GR) term makes use of the weight matrix \mathbf{W} and imposes smoothness of the function over the graph. The second term is dependent on the known label information and measures deviations from *model implied* label information. Consequently, this term is often called the *data fitting* term. In the traditional transductive setting, labels for a subset of nodes, \mathbf{S} are given and the remaining labels need to be inferred. For example, in the GFHF method, the

function \mathbf{F} is set to \mathbf{Y} for the labeled nodes \mathbf{S} and is estimated for the unlabeled nodes by minimizing only the graph regularization term $Q_{GR}(\mathbf{F}) = \sum_{k=1}^K \mathbf{F}_k^T \mathbf{L}_{un} \mathbf{F}_k$. In the LGC method, the function \mathbf{F} is estimated by minimizing $C \sum_{k=1}^K \mathbf{F}_k^T \mathbf{L}_{nrm} \mathbf{F}_k + \|\mathbf{F}_k - \mathbf{Y}_k\|^2$ where \mathbf{Y} is set to zero for the unlabeled nodes.

As explained in section 2, applying the methods in solution setting 1 is straightforward. The selection of S will be discussed below. Some care is needed when extending the methods to deal with solution setting 2, which aims to make effective use of the class distribution information \mathbf{P} . It makes good sense to set $\mathbf{Y} = \mathbf{P}$ in (3.1) and thus try and force \mathbf{F} to be close to \mathbf{P} while also ensuring its smoothness on the graph. It is also a good idea to try a bit more and use the uncertainty present in the class distribution information to apply different weights to different data fitting terms. With this in mind we give a generic quadratic formulation with two parameters (\mathbf{H} and Λ). Original versions of the GFHF method and the LGC method follow as special cases. Also, choosing the parameters differently using \mathbf{P} leads to extensions of these methods to solution setting 2.

3.1 Generic Quadratic Objective Function. We define the following generic quadratic objective function to estimate \mathbf{F} and we shall see that the objective functions used in GFHF and LGC methods are special cases of this objective function.

$$(3.2) \quad Q_g(\mathbf{F}) = \sum_{k=1}^K C(\mathbf{F}_k - \mathbf{V}\Lambda\mathbf{Y}_k)^T \mathbf{H}(\mathbf{F}_k - \mathbf{V}\Lambda\mathbf{Y}_k) + \mathbf{F}_k^T \mathbf{L} \mathbf{F}_k$$

where $\mathbf{F} = [\mathbf{F}_1 \cdots \mathbf{F}_K]$ and \mathbf{Y}_k is k th column of \mathbf{Y} . On comparing (3.2) and (3.1) we see that the data fitting term is nothing but a weighted quadratic error function and the graph regularization term is another quadratic function defined in terms of graph Laplacian matrix \mathbf{L} .

Let us consider the weighted quadratic error function. We have introduced two parameters: a generic error weight matrix \mathbf{H} and a label degree matrix Λ . The role of \mathbf{H} is to give different weightage to individual errors (data fitting terms) and is assumed to be positive definite. When \mathbf{H} is diagonal the data fitting term is essentially a weighted sum of squared errors with error in the i th node weighted by h_{ii} . Λ is a diagonal matrix and its role is to incorporate any label degree information we want to associate with each node where $0 \leq \lambda_{ii} \leq 1$. For example if node i is unlabeled then $\lambda_{ii} = 0$ and when it is labeled *fully* $\lambda_{ii} = 1$. Note that the interval $0 < \lambda_{ii} < 1$ brings in the notion of *partial* labeling and provides flexibility when \mathbf{y}_i (i th row of \mathbf{Y}) is inaccurate.

Finally, \mathbf{V} is a node regularization matrix which

is also diagonal. Wang et al[7] introduced the notion of node regularization in the original LGC formulation to handle *label imbalance* problem in the graph. Assuming \mathbf{y}_i to contain only one non-zero element they defined normalized label matrix $\hat{\mathbf{Z}} = \hat{\mathbf{V}}\mathbf{Y}$ with $v_{ii} = \sum_{k=1}^K \frac{y_{i,k} d_{ii}}{\tilde{\eta}_k}$. Here $\tilde{\eta}_k = \sum_{i=1}^n d_{ii} y_{i,k}$ and we can see that $\sum_{i=1}^n \tilde{z}_{i,k} = 1$. Thus \mathbf{V} balances the influence of labels from different classes and allows with high degree to make more contribution. In our general setting we allow $0 \leq y_{i,k} \leq 1$, $k = 1, \dots, K$ subject to $\sum_{i=1}^K y_{i,k} = 1$ and we also have the label degree matrix Λ . Therefore we redefine v_{ii} as: $v_{ii} = \sum_{k=1}^K \frac{d_{ii} \lambda_{ii} y_{i,k}}{\eta_k}$ where $\eta_k = \sum_{i=1}^n d_{ii} \lambda_{ii} y_{i,k}$. Defining $\mathbf{Z} = \mathbf{V}\Lambda\mathbf{Y}$ we have $\sum_{i=1}^n z_{i,k} = 1$, $\forall k$ and, \mathbf{Z} becomes a normalized matrix as earlier.

Given the additive nature of the function $\mathbf{Q}_g(\mathbf{F})$, \mathbf{F}_k can be solved independently. Then setting the derivative $\frac{\partial \mathbf{Q}_g(\mathbf{F})}{\partial \mathbf{F}_k}$ to zero gives the following closed form solution that involves matrix inversion and costs $O(n^3)$.

$$(3.3) \quad \hat{\mathbf{F}}_k = C(\mathbf{L} + C\mathbf{H})^{-1} \mathbf{H}\mathbf{Z}_k$$

Here \mathbf{Z}_k is k th column of \mathbf{Z} . In the following discussion we assume that \mathbf{H} is diagonal and is invertible. Then equation (3.3) can be rewritten as: $\hat{\mathbf{F}}_k = (\mathbf{I} + \frac{1}{c}\mathbf{H}^{-1}\mathbf{L})^{-1} \mathbf{Z}_k$. Then instead of estimating $\hat{\mathbf{F}}_k$ from (3.3) using matrix inversion, we can find it by iteratively using the fixed point equation: $\hat{\mathbf{F}} = -\frac{1}{c}\mathbf{H}^{-1}\mathbf{L}\hat{\mathbf{F}} + \mathbf{Z}$. Note that the convergence rate of the iterative solution depends on the eigen values of the matrix $(\mathbf{I} + \frac{1}{c}\mathbf{H}^{-1}\mathbf{L})$. Also the fixed point equation changes depending on whether normalized or unnormalized graph Laplacian is used and the specific form of error weight matrix \mathbf{H} .

3.2 GFHF Method As a Special Case. The GFHF method [9] was originally proposed to solve semi-supervised learning problem based on a Gaussian random field model for a weighted graph model. In this formulation a real valued function \mathbf{f} is estimated by minimizing a quadratic energy function ($\sum_{k=1}^K \mathbf{F}_k^T \mathbf{L}_{un} \mathbf{F}_k$) subject to the constraint that the function \mathbf{f} takes the actual label values on the set of labeled nodes. It is assumed that $y_{i,k}$ takes value in $\{0, 1\}$. This solution is retrieved by considering our generic quadratic formulation and setting $\mathbf{L} = \mathbf{L}_{un}$, $\mathbf{V} = \mathbf{I}$, $\mathbf{H} = \mathbf{D}\Lambda(\mathbf{I} - \Lambda)^{-1}$ and $C = 1$. Further $\lambda_{ii} = 1$ if node i is labeled and is zero otherwise. Note that the specific structure of \mathbf{H} and Λ imposes the constraints on the labeled nodes by assigning infinite weights to the errors on labeled nodes. On substituting these specific matrices we get:

$$(3.4) \quad \mathbf{F}_k = (\mathbf{I} - (\mathbf{I} - \Lambda)\mathbf{D}^{-1}\mathbf{W})^{-1} \Lambda\mathbf{Y}_k$$

where \mathbf{Y}_k is k th column of the matrix \mathbf{Y} . The corresponding fixed point equation (after including the node regularization matrix \mathbf{V}) is: $\mathbf{F}_k = (\mathbf{I} - \Lambda)\mathbf{D}^{-1}\mathbf{W}\mathbf{F}_k + \Lambda\mathbf{Z}_k$.

Note that (3.4) can handle both the solution settings presented in section 2. In a transductive setting with labeled and unlabeled nodes we have $\lambda_{ii} = 1$ for the *labeled* subset of nodes such that $i \in \mathbf{S}$ and $\lambda_{ii} = 0$ for the remaining set of nodes $i \in \bar{\mathbf{S}}$. [9] also gave a random walk interpretation to this method. That is, $\hat{f}_{i,j}$ has an interpretation as the probability that a particle starting from node i , hits a labeled node with label 1 (with class j being considered as 1 like in one-vs-all setting). They also discussed incorporating external class label information using the notion of dongle nodes. Here each node having label information is attached with a dongle node and the label value is tied to this node. Further a transition of probability μ was defined from each node to its dongle node and all other transitions are covered with probability $(1 - \mu)$.

It may be noted that (3.4) is a general form with this interpretation and when $\Lambda = \mu\mathbf{I}$ we get the original form of GFHF method. Thus we allow these transition probabilities to take different values for different nodes and this is very useful since we may have different levels of confidence on the label information we get for each node from the external classifier. Below we shall discuss how to choose these values.

3.3 LGC Method As a Special Case and its Extension. Zhou et al[8] proposed a semi-supervised learning approach to design a *smooth* classification function \mathbf{F} such that it respects any intrinsic structure revealed by the labeled and unlabeled nodes. Starting with the intuition of label propagation they introduced an iterative algorithm which essentially finds the solution to a FP equation. With \mathbf{y}_i having only one non-zero element and by setting, in our generic quadratic formulation, $\mathbf{L} = \mathbf{L}_{nm}$, $\mathbf{V} = \mathbf{I}$, $\mathbf{H} = \mathbf{I}$, $\lambda_{ii} = 1$, $\forall i \in \mathbf{S}$ and $\lambda_{ii} = 0$, $\forall i \in \bar{\mathbf{S}}$ we can retrieve the FP equation following the same steps in the previous section. Zhou et al[8] also considered several variants of fixed point equations and showed how one such FP equation can be obtained as solution from optimizing a regularized objective function ($C \sum_{k=1}^K \mathbf{F}_k^T \mathbf{L}_{nm} \mathbf{F}_k + \|\mathbf{F}_k - \mathbf{Y}_k\|^2$). Equation (3.2) can be seen as a generalized version of such an objective function. Unlike GFHF method the regularization constant C is to be set using standard techniques like cross-validation and this is possible only when sufficient number of labeled nodes is available.

Extension for Setting 2 We extend the basic LGC method to handle setting 2 as follows. Firstly, we relax \mathbf{y}_i as described earlier. Secondly, we make use

of the label degree matrix Λ to handle inaccurate label information. Finally, we allow \mathbf{S} to be the entire set of nodes. With specific structures of \mathbf{L} and \mathbf{H} defined as above we get:

$$(3.5) \quad \mathbf{F}_k = (1 - \gamma)(\mathbf{I} - \gamma\mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}})^{-1}\Lambda\mathbf{Z}_k$$

where $\gamma = \frac{1}{1+C}$ and it can re-written in fixed point form as done earlier with the GFHF method. We will discuss shortly how to choose C and Λ in the two settings.

The solution of the fixed point equations, (3.4) and (3.5) can be obtained iteratively. In all our experiments we found the solution iteratively. The iterative update is very useful when \mathbf{W} is sparse (which is often the case in practice) and the convergence rate is dependent on the eigen values of the matrix involved in the inversion. In the worst case the cost is expected to be in the order of $O(mn^2)$ where m is a bound on the number of neighbors of a node in the graph. When the graph is dense, the complexity can be cubic in n [1].

3.4 Parameter Selection in Two Settings. In this section we present ways of selecting the subset \mathbf{S} and other parameters Λ and C that appear in (3.2). We propose two scoring schemes that can be used to make the subset selection for solution setting 1.

Maximum Probability Scoring Scheme Let $p_i^{max} = \max_l p_{i,l} \forall i$. Then we sort the nodes based on the maximum probability score (p_i^{max}) of each node from its class distribution. Then the subset of nodes that satisfy the condition $p_i^{max} \geq p_{th}$ can be selected as the set of labeled nodes. Essentially we select nodes on which we have high confidence in their labeling assignments. If p_{th} is very high then the number of labeled nodes becomes smaller. However the noise level (percentage of noisy labeled nodes) will also be low and care has to be taken to ensure that each class has at least some labeled nodes. On the other hand, if the threshold level is low then though the number of labeled nodes increases the noise level also increases. Alternately, we can also select top-M percentage of nodes as the subset of labeled nodes. We will refer to the subset selection scheme based on the maximum probability score as *MPS* scheme.

Entropy based Scoring Scheme In this scheme we sort the nodes based on an entropy based score η_i for node i defined as: $\eta_i = 1 - E(\mathbf{p}_i)/E_{max}$ where $E(\mathbf{p}_i)$ represents the entropy of the class distribution \mathbf{p}_i and $E_{max} = \log(K)$ is the maximum entropy possible for a given number of classes K . Thus η_i lies in the interval $[0, 1]$ and takes high or low values depending on the spread of the class distribution. This score is motivated from the view that we do not have any class label information when the class distribution is uniform and

we are certain when $\mathbf{p}_{i,k} = \delta_{k,c_i}$. Setting a threshold to select the subset of labeled nodes \mathbf{S} is harder here compared to *MPS* scheme and as mentioned earlier, so we can simply select the top-M percentage of nodes sorted by this score. We will refer to this subset selection scheme as *EBS* scheme.

Note that the above two schemes only use the class distribution information to select the subset of nodes. It is also possible in addition to make use of the graph structure information and is a direction for further work.

Choice of Λ Recall that λ_{ii} has the interpretation of transition probability in GFHF method; and it has the interpretation of *partial* label information in the LGC-extension method. Earlier we gave two useful scoring schemes based *only* on the external classifier information. We can make use of either of these schemes to fix λ_{ii} . That is, we can either set $\lambda_{ii} = p_i^{max}$ or $\lambda_{ii} = \eta_i$. We evaluate both these schemes in our experiments. Note that when the subset selection scoring scheme also makes use of graph structure information, then we may not want to use the same scoring scheme to choose λ . This is because we may like to give importance to the label degree or transition probability to each node *only* based on the information from the external classifier.

Choice of C We use 5-fold cross-validation (CV) technique to choose C for LGC. In the first setting, for each value of C chosen from a range in the log scale, we evaluate 5-fold CV accuracy on the selected subset of nodes using the label information available for these nodes. Then we pick the value of C that gives the maximum accuracy and find the final solution using the chosen C value. Here the label information is nothing but the label that has the maximum probability score. In the second setting we evaluate 5-fold CV accuracy on top-M percentage of nodes selected using *MPS* or *EBS* scheme. It is worth pointing out that, inaccuracies in the external class distribution could lead to a choice of C that is tilted away from the best possible choice.

4 Probabilistic Methods

The methods in this class, viz., the information regularization (IR), dual IR methods and the probabilistic weighted vote relational neighbor (WvRN) classifier method, estimate the class distributions of the nodes. In this section we discuss their adaptation for the two solution settings and also introduce least squares regularization (LSR) based method. Overall, we group the IR, DIR and LSR methods under the category of region based regularization methods.

4.1 WvRN Method and its Extension. The original probabilistic weighted vote relational classifier (with relaxation labeling) method [4] was formulated

to solve the collective classification problem where class distributions of a subset of nodes are known and fixed. Then the class distributions of the remaining (unlabeled) nodes are obtained by an iterative algorithm. We give a version of this algorithm in Algorithm 4.1. It has two components, namely, weighted vote relational neighbor classifier (WvRN) component and relaxation labeling (RL) component. The relaxation labeling component performs collective inferencing and keeps track of the current probability estimates $\mathbf{p}_i^{(t)}$ for all unlabeled nodes at each time instant t . These *frozen* estimates $\mathbf{p}_i^{(t)}$ are used by the relational classifier. The relational classifier computes the probability distribution for each unlabeled node as the weighted sum of probability distributions $\mathbf{p}_j^{(t)}$ of its neighbors with weight w_{ij} . Since relaxation labeling may not converge, simulated annealing is performed to ensure convergence (as given in step (2) of the algorithm) [4]. Note that $\beta^{(t)} \rightarrow 0$ for sufficiently large number of iterations; therefore, convergence is guaranteed. It has been observed that the performance is robust when $0.9 < \nu < 1$. The WvRN algorithm (Algorithm 4.1) can be directly used to solve the problem in the first setting by setting the probability distributions associated with the selected set of nodes $i \in \mathbf{S}$ to δ_{k,c_i} and initializing the probability distributions of the remaining set of nodes ($\bar{\mathbf{S}}$) with the class prior obtained using the label information of the selected set of nodes. Here the label for a node (c_i) in the selected subset is defined as the label with the maximum probability score obtained from p_i .

Extensions for Setting 2 Next we extend the WvRN algorithm to solve the problem in setting 2. We consider two variants. In the first variant we consider a simpler form where we initialize $\mathbf{p}_i^{(0)}$ with the external classifier information for all nodes and run Algorithm 4.1 as it is. In the second variant we use the dongle node idea (used in the GFHF method) and modify the relational classifier estimate from \mathbf{q}_i to $\tilde{\mathbf{q}}_i$ as follows. With λ_{ii} representing the transition probability, we define $\tilde{q}_{i,k} = \lambda_{ii}p_{i,k}^{(0)} + (1 - \lambda_{ii})q_{i,k}$ (where $q_{i,k}$ is as defined in Algorithm 4.1) and $\mathbf{p}_i^{(0)}$ is the distribution information available from the external classifier. We can select Λ using MPS or EBS schemes, and ν using the cross-validation technique (see section 3.4).

4.2 Region Based Regularization Methods Corduneanu and Jakkola [2] proposed the information regularization method. They introduced a notion of region, where each region is defined as a subset of nodes in the graph. The intuition is that nodes belonging to a given region have the same label. Here a weight is associated with each region and a weight to each node that

Algorithm 4.1 Probabilistic WvRN Algorithm

- Set $t = 0$, $\beta^{(0)} = 1$ and $\nu=0.95$
 - For all nodes $i \in \bar{\mathbf{S}}$, initialize $\mathbf{p}_i^{(0)}$ to class prior (obtained from known labeled nodes)
 - Until convergence holds for all the nodes in $\bar{\mathbf{S}}$ do:
 - For each element $i \in \bar{\mathbf{S}}$ and $k = \{1, \dots, K\}$
 1. Estimate node class probability (using neighbor information)

$$q_{i,k} = \frac{1}{\psi} \sum_j w_{i,j} p_{j,k}^{(t)}$$
 (where ψ is a normalizing constant)
 2. Set $p_{i,k}^{(t+1)} = \beta^{(t)} q_{i,k} + (1 - \beta^{(t)}) p_{i,k}^{(t)}$
 - Set $t = t + 1$ and $\beta^{(t+1)} = \beta^{(t)} * \nu$
-

defines the relative importance of points that belong to a given region. Further, a distribution is associated with each region and node. Then the distributions of labels are propagated on a graph for semi-supervised or transductive learning. We refer to methods that work within this framework of region as region based regularization methods. The IR, DIR and LSR methods fall in this category. In graph based classification problem, each edge forms a region and the weight of the region is nothing but the edge weight; further, an equal degree of importance is given to both the nodes connected by the edge. Then the regularized optimization problem can be written as minimization of the objective function:

$$(4.6) \quad H(\mathbf{P}; \mathbf{W}) = C \sum_{i \in \mathbf{S}} \lambda_{ii} D(\mathbf{p}_i^{(0)}, \mathbf{p}_i) + \sum_{i,j} w_{i,j} D(\mathbf{p}_j, \bar{\mathbf{p}}_{i,j})$$

where $D(\cdot)$ denotes a divergence measure, λ_{ii} denotes *relative* weight factor that we would like to assign to i -th node in \mathbf{S} . Further, $\bar{\mathbf{p}}_{i,j}$ denote the probability distribution associated with the region (edge) (i, j) . The optimization problem should include the following constraints: (1) $0 \leq p_{i,k} \leq 1$ and (2) $\sum_{k=1}^K p_{i,k} = 1$. The region distributions are constrained in a similar way. In this formulation, an alternating learning algorithm[2] consists of two steps. In the first step, *all* the node distributions $\mathbf{p}_i, \forall i$ collected as a vector \mathbf{P} are fixed and, *all* the region distributions $\bar{\mathbf{p}}_{i,j}, \forall (i, j)$ pairs are obtained by minimizing *only* the second term in (4.6). In the second step, using these estimated region distributions, the node distributions are re-estimated by minimizing *both* the terms in (4.6). These steps are repeated until convergence.

Now let us look at how (4.6) is used in the two solution settings. In the first setting, having selected the subset of nodes \mathbf{S} we fix $p_{i,k} = p_{i,k}^{(0)} = \delta_{k,c_i}, \forall i \in \mathbf{S}$

and $k = 1, \dots, K$. Then the optimization is over $\mathbf{P}_{\mathbf{V} \setminus \mathbf{S}}$ where we have explicitly indicated the set of nodes to be optimized. Therefore the first term in (4.6) does not play a role in the optimization process. In the second setting, the first term also plays a role since we optimize over the entire set of node distributions $\mathbf{P}_{\mathbf{V}}$. Using this broad setup let us consider several methods that stem from different divergent measures $D(\cdot)$, namely, KL-divergence and squared error (loss). The complexity of all these methods is same as that of the function estimation methods.

4.2.1 Information Regularization (IR) Method

Within the above framework, Corduneanu and Jakkola [2] used KL-divergence as the divergence measure $D(\cdot)$ and called the second term as *information regularization* term. Then closed form solution can be obtained for $\bar{\mathbf{p}}_{i,j}$ and is given by:

$$(4.7) \quad \bar{\mathbf{p}}_{i,j} = \frac{1}{2}(\mathbf{p}_i + \mathbf{p}_j).$$

Note that $\bar{\mathbf{p}}_{i,j}$ s are the region (edge) distributions and are obtained from minimizing the second term in (4.6) with the KL-divergence measure. As mentioned earlier, in solution setting 1 the first term does not play a role in the optimization process. Then the distributions for the unlabeled nodes are directly obtained as:

$$(4.8) \quad p_{i,k} = \frac{1}{\chi_i} \exp\left(\sum_{i,j} w_{i,j} \log \bar{p}_{i,j}(k)\right)$$

where $\bar{p}_{i,j}(k)$ denotes k th component in $\bar{\mathbf{p}}_{i,j}$ (see (4.7)) and χ_i denotes a normalizing constant. Thus in setting 1, closed form solutions exist in *both* the steps of the alternating learning algorithm. In solution setting 2, with the first term also playing a role in the optimization process there is *no closed form solution* like (4.8). Then constrained optimization using either Newton's method or exponentiated gradient algorithm is carried out. This step could be expensive and affects the scalability of the method. To address this issue, Tsuda [6] proposed a dual information regularization method, where closed form solutions are obtained in both the steps *even* when the first term in (4.6) plays a role.

4.2.2 Dual Information Regularization (DIR) Method

Tsuda [6] modified the regularizer by interchanging the arguments in the KL divergence measure ($KL(\bar{\mathbf{p}}_{i,j} \parallel \mathbf{p}_i)$) and using modified $\bar{\mathbf{p}}_{i,j}$ given by:

$$(4.9) \quad \bar{\mathbf{p}}_{i,j} = \frac{1}{Z_{ij}} \exp\left(\frac{1}{2}(\log(\mathbf{p}_i) + \log(\mathbf{p}_j))\right)$$

where Z_{ij} is the normalization factor. Then the closed form solutions for \mathbf{p}_i are obtained as:

$$(4.10) \quad \mathbf{p}_i = \frac{1}{\lambda_{ii} + D_{ii}} (\lambda_{ii} \mathbf{p}_i^{(0)} + \sum_{i,j} w_{i,j} \bar{\mathbf{p}}_{i,j})$$

where $D_{ii} = \sum_j w_{i,j}$. We refer to the update equations (4.9) and (4.10) as the dual information regularization (DIR) method. Let us consider the implications in the two settings. In the first setting we fix the node distributions of the nodes in \mathbf{S} and estimate only the node distributions of the nodes $\mathbf{V} \setminus \mathbf{S}$. This estimation is done using (4.9) and with $\lambda_{ii} = 0, \forall i \in \mathbf{S}$ in (4.10). In the second setting, (4.9) and (4.10) are used as they are. Thus, closed form solutions are available in both settings and in both steps of the learning algorithm and, this helps in improving the speed of the IR method.

4.2.3 Least Squares Regularization (LSR) Method

We propose to use the squared error as the divergence measure $D(\cdot)$; that is, for any two distributions \mathbf{p} and \mathbf{q} , we define $D(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|^2$. Then it is easy to verify from (4.6) that the optimal $\bar{\mathbf{p}}_{i,j} = \frac{1}{2}(\mathbf{p}_i + \mathbf{p}_j)$ and is *same* as the one obtained in the IR method. In this method, we proceed as in the 2-step algorithm where we estimate $\mathbf{P}_{\mathbf{V}}$ keeping all the node distributions fixed. This results in the closed form solution: $\mathbf{p}_i = \frac{1}{\lambda_{ii} + D_{ii}} (\lambda_{ii} \mathbf{p}_i^{(0)} + \sum_{i,j} w_{i,j} \bar{\mathbf{p}}_{i,j})$. This solution is *exactly same* as the one given in (4.10). We refer the updates (4.7) and (4.10) as the least squares regularization (LSR) method. Note that application of this method in two settings is exactly same as described above for the DIR method.

Relation to other Methods The LSR objective function, that is, (4.6) with squared error as the divergence measure has structure similar to the generalized quadratic function (3.2). The key difference is that unlike the rows of \mathbf{F} , the node distributions \mathbf{p}_i and $\mathbf{p}_i^{(0)}$ are constrained to be probability distributions. Now, comparing the DIR and LSR updates, we see that they differ in the way region distributions are updated. *Thus, the LSR method interestingly combines the IR based region distribution update with the D-IR based node distribution update. Also, like the IR methods, the objective function is convex and has global minimum.*

5 Experiments

We conducted two experiments on the two function estimation methods, LGC and GFHF, and the four probabilistic methods, WvRN, IR, DIR and LSR. In different settings the acronyms of methods will appropriately refer to the modifications that we described earlier in the paper. For WvRN in solution setting 2 there are

Table 1: Public Datasets Description. n , K , $|E|$ and d denote the number of nodes, classes, edges and dimensions respectively.

Dataset	n	K	$ E $	d
WebKB	1051	2	269046	4840
USPS	3874	4	6398	256
CoraCite	4270	7	22516	-
CoraAll	4270	7	71824	-

two versions, which we will refer to as WvRN-V1 and WvRN-V2 (see section 4.1 and 4.2). In the first experiment we evaluate the performance of the methods in solution settings 1 and 2 with four publicly available benchmark datasets. Next, we present results from the second experiment where we evaluated the performance in solution setting 2 on datasets constructed from web pages of three shopping sites.

5.1 Datasets Description. The details of the datasets are given in Table 1. The WebKB dataset contains two document categories, *course* and *noncourse*. Each document has two types of information, the webpage text content and link information. The number of features in page and link representations are 3000 and 1840 respectively. Following [7] we constructed the graph based on cosine distance neighbors with Gaussian weights and chose 200 nearest neighbors. The USPS dataset is a handwritten digit recognition task, for which we used the same setting as given in [8]. The number of features is 256, obtained from a 16×16 image. The four classes correspond to digits 1 to 4. The graph is constructed using a radial basis function kernel with width set to 1.25; the number of neighbors is set to 1. The CORA dataset comprises computer science research papers. There are seven classes associated with the papers; the classes correspond to the following machine learning sub-topics: *Case-based Methods*, *Genetic Algorithms*, *Neural Networks*, *Probabilistic Methods*, *Reinforcement Learning*, *Rule Learning* and *Theory*. The dataset consists of the full citation graph with labels for the topic of each paper. There are two variants of this dataset, referred to as CORACITE and CORAALL. These variants come from the way the papers are linked. The CoraCite variant uses *only* citation link and an edge is placed between two papers if one cites the other. The weight of an edge is normally one unless the two papers cite each other, in which case it is two. The CoraAll variant uses *both* citation and author link information, where an edge is placed (in addition to co-citation) when two papers have author relation as well. We used the CORA dataset versions available

from Netkit package described in [4].

5.2 Classifier Information Generation. While these datasets have only label information we need inaccurate external classifier (distribution) information in our problem formulation. Therefore we need a model which generates this distribution given the label information. We now describe the model used in our experiments. We fix two probability parameters p_{min} and p_{max} with $p_{min} < p_{max}$ that take values in $[0, 1]$. Given p_{min} and p_{max} we generate the distribution for each node as follows. In the first step generate a random number p_{label} from the interval $[p_{min}, p_{max}]$ and treat p_{c_i} as the probability score of the true label (c_i). Then we generate $K - 1$ random numbers (pr_k with $k \neq c_i$) from the interval $[0, 1]$ and assign $p_k = \frac{pr_k(1-p_{c_i})}{\psi}$ where ψ is a normalizing constant such that $\sum_{k \neq c_i} p_k = 1 - p_{c_i}$. The choices of p_{min} and p_{max} determine the degree of inaccuracy present in the information. If p_{min} is set too low many nodes get labeled wrongly. Also note that the number of classes play a role in determining the degree of accuracy since the mass $1 - p_{c_i}$ gets distributed across $K - 1$ number of classes. We considered different levels of inaccuracy by setting different values for p_{min} and p_{max} . As the observations were almost the same across the methods and settings for different levels of inaccuracy, we present results only for one set of values.

5.3 Experiment With Setting 1. We conducted the experiment with the two function estimation methods, LGC and GFHF, and the four probabilistic methods, WvRN, IR, DIR and LSR. We measured the classifier accuracy on the entire set of nodes. This experiment was conducted with both maximum probability scoring (MPS) and entropy based scoring (EBS) subset selection schemes. The selected subset size $|S|$ was varied from *top* 10 percent to 90 percent (as per the chosen selection scheme).

Let us discuss parameter selection. Recall that in the GFHF method C is set to 1. In the IR, DIR and LSR methods only the second term is present (once we fix the probability distribution of *labeled* nodes). In the WvRN method we set $\nu = 0.95$. Thus only the LGC method required tuning of the regularization parameter C . For selecting C we used 5-fold CV on the labeled nodes and varied C in the range $[0.00153, 100]$ (doubling in each step). As mentioned in section 3.4.3, note that the 5-fold evaluation used here is inaccurate since it is based on noisy information. Therefore, the C estimate can be inferior sometimes.

The average classifier accuracy was computed as the average of accuracies obtained from 100 random partitions of the graph. The results for one parameter

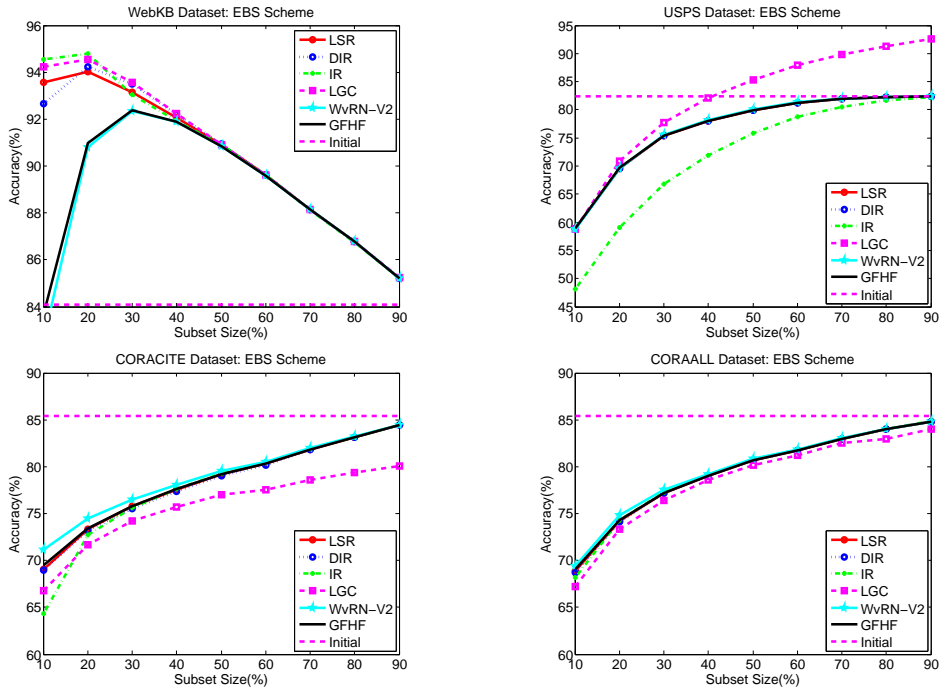


Figure 1: Setting 1 Results: X-axis represents the subset size in terms of percentage of number of nodes and Y-axis represents accuracy. The results correspond to the EBS subset selection scheme. Results for the MPS subset selection scheme are not given because of their closeness with EBS; see text. *Initial* refers to accuracy computed from generated inaccurate distribution.

setting with the MPS subset selection scheme for all the datasets are given in Figure 1; we set $p_{max} = 0.99$ for all the datasets and set p_{min} as 0.4, 0.2 and 0.1 for WebKB, USPS and CORA datasets (both CORACITE and CORAALL) respectively.

From Figure 1 we see that only on the WebKB dataset all the methods are able to perform better than the initial accuracy; in fact, even on this dataset, WvRN and GFHF methods fall short when the subset size $|S|$ is small. All methods other than LGC give almost the same accuracy as $|S|$ becomes larger, on all datasets; of course, the exact value of $|S|$ at which they converge vary from one dataset to another. On the CORA datasets WvRN performed better followed by GFHF, LSR, DIR, IR and LGC when $|S|$ is small. On comparing EBS and MPS selection schemes we found that the EBS scheme performed slightly better (around 2%) on CORA datasets, while the performances were almost same on WebKB dataset; the EBS scheme performed slightly inferior on USPS dataset. Except for these variations, the behavior of the EBS scheme as a function of $|S|$ was almost same as that of the MPS scheme; therefore, only the results for the EBS scheme are shown in the figure. We analyzed the inferior performance of LGC on CORA datasets and observed that incorrect choice of C due to noisy CV estimates was the reason behind it.

5.4 Experiment With Setting 2. We conducted the experiment with LGC, GFHF, IR, DIR, LSR, WvRN-V1 and WvRN-V2. The classifier accuracy measurement, evaluation with two selection schemes, subset size variation and p_{min} and p_{max} settings remain the same as in setting 1. In this setting IR, DIR, LSR and WvRN-V2 *also* require parameter tuning (that is, C and ν). In the case of WvRN-V1 there is no parameter tuning and we set $\nu = 0.95$. In the case of GFHF we set $C = 1$ as done earlier. All parameter tunings are done using 5-fold CV, using top M percent of selected nodes, as in setting 1. It is useful to recall that the main difference of setting 2 from setting 1 is that the distribution information from all nodes is used during the solution process. We used the same range for C as in setting 1 for the LGC method. In the case of WvRN-V2 method we used the same range but set $\nu = \frac{1}{1+C}$. In the modified IR method we selected C from the range $[0.0625, 312.5]$ (doubling in each step). In the case of LSR and DIR methods we selected C from the range $[0.078, 10]$ (doubling in each step). Finally the average accuracy was computed from 100 random partitions of the graph for all the methods except IR.

Efficiency The objective function optimization in the case of IR takes significantly longer time (an order of magnitude) compared to other methods in this setting due to the nonlinear optimization involved. Therefore, only for IR, we computed the average performance from

100, 20 and 10 partitions for USPS, WebKB and CORA datasets respectively. In terms of speed, WvRN, LGC, GFHF, LSR, DIR and IR methods come in that order. Although the update equations for the DIR and LSR methods look similar, the DIR method took more time to converge than the LSR method and we observed that the LSR method was approximately 6-10 times faster than the DIR method. On the other hand, the LSR method was slower by 2-4 times compared to the remaining methods.

Choice of C from Cross-Validation The choice of C from the noisy CV accuracy estimates has an important effect on the classification performances of LGC, IR, DIR and LSR methods; this effect varies across datasets. In these methods we observed that inferior performances were obtained when we chose C that gave the best CV accuracy. To get improved and robust performances, we recommend an alternate way to choose C : choose the smallest C such that the corresponding CV accuracy estimate was within, say, 5% of the best CV accuracy. This helps because the variation of CV accuracy estimate as a function of C can be quite flat around the best CV accuracy and, in such cases choosing the least C within a specified accuracy estimate regularizes the solution better. The results reported in Figure 2 are the improved performances obtained using the recommended way of choosing C described above; some inferior performances can be still seen with IR on the WebKB, USPS datasets (particularly for low values of $|\mathbf{S}|$) and LGC on CORA datasets. From this viewpoint, the results suggest that DIR, LSR and WvRN-V2 are more robust compared to IR and LGC.

Classifier Performance From figure 2 we clearly see that WvRN-V1, WvRN-V2, IR, DIR and LSR methods improved the performance significantly over the initial accuracy on all datasets. The LGC-extension and modified GFHF methods improved the performance significantly on WebKB and USPS datasets. The performance curves of GFHF and WvRN-V1 remain flat in each plot. This is because there is no parameter tuning involved in these methods. For the other methods the choice of C made remained almost same for all $|\mathbf{S}|$. Considering the performance on all the datasets the EBS scheme seems to be more robust compared to the MPS scheme, although slightly inferior performance is seen on the CORAALL dataset. Note that the performance variations over $|\mathbf{S}|$ is lesser with the EBS scheme particularly on the WebKB dataset. We believe this behavior is due to the conservative estimate of λ_{ii} given by this scheme compared to the MPS scheme in the noisy scenario. Clearly, compared to the first setting, the second setting using distribution information of all nodes

during the solution process enhances the performance significantly. Though function estimation methods performed quite close on two datasets, the probabilistic methods performed better considering all the datasets. Among the probabilistic methods, the performances of IR, DIR and LSR were quite close as $|\mathbf{S}|$ became large. Comparing WvRN-V2 and WvRN-V1, since the performance difference is significant ($>4\%$ in many cases), WvRN-V2 is to be preferred. Since the performance variation across C in the WvRN-V2 method is not much it seems that much of the gain comes from using dangling nodes. Recall that GFHF and WvRN methods have almost same performances on all datasets in the first setting; however, in setting 2, distinctly different performances of these methods are seen, particularly on CORA datasets. More investigation is needed to understand this behavior. Finally, the performance improvement is dependent on the quality of relational graph (with respect to the assumption of strong connectivity of nodes belonging to same class) and initial accuracy.

5.5 Experiment with Shopping Datasets. We evaluated the performances of DIR, LGC, WvRN-V2 and LSR in solution setting 2 on real world datasets from three shopping sites (*www.compusa.com*, *www.uncommongoods.com* and *www.walmart.com*). We considered two binary classification problems: *product detail vs non-product* and *product listing vs non-listing*. While the product pages are about one specific product like canon camera with certain model number in more detail, the listing pages are about several products of same category (for example, different models of canon camera) arranged as a list in each page. These problems along with their site names are referred to as CU-D, CU-L, UG-D, UG-L, WM-D and WM-L respectively in Table 2. In each problem, the class distribution score for each page was obtained using an external classifier (EC) based on content features. The relational graph was constructed using structural signature (shingle) obtained using html tags of web pages. An edge between two pages was formed when their structural signatures had a match score of at least 6 (the values are in the range $[0,8]$) and a unit weight was assigned to each such edge. Also, each node was connected to a maximum of 20 other nodes. A subset of pages (nodes) in each site (graph) was manually labeled; the number of labeled nodes is indicated as L in Table 2. The classifier accuracy was evaluated for the external classifier and the various methods on these labeled nodes. The results are given in Table 2.

In almost all the cases the EBS scheme performed better compared to the MPS scheme. The results given in Table 2 are the best performances obtained over the

Table 2: Experiment: Shopping Datasets. n , L and $|E|$ denote the number of nodes, labeled nodes and edges in the graph respectively. The classification accuracy evaluated on the labeled nodes for the external classifier (EC), LGC, WvRN-V2, LSR and DIR methods are given.

Problem	n	L	$ E $	EC	LGC-Ext	WvRN-V2	LSR	DIR
CU-D	53023	1433	1186190	77.18	77.88	77.74	78.02	78.02
CU-L	53023	1433	1186190	91.70	93.79	93.93	94.21	94.21
UG-D	82027	1166	1714228	94.51	97.68	94.51	95.54	95.80
UG-L	82027	1166	1714228	81.90	91.51	85.93	91.25	91.51
WM-D	67997	1250	1318903	93.20	94.48	95.12	95.52	95.60
WM-L	67997	1250	1318903	93.36	94.72	94.72	95.60	95.04

percentage of subset sizes ($|\mathcal{S}|$) and C values. There were minor variations (0.5 – 2%) in the performances over ($|\mathcal{S}|$) and choices of C . The results clearly demonstrate that significant performance improvement (as high as 10%) can be achieved. Also, the usefulness of the proposed extensions and LSR method as effective alternate methods can be seen from comparison with the DIR method.

6 Summary

We considered the problem of collectively classifying entities where relational information and inaccurate class distribution information from another (external) classifier are available. We present below a list of key observations from the experimental studies conducted on several benchmark and real world datasets.

- Of the two solution settings evaluated, the second setting (which uses external classifier information of all nodes) is better. Using this second setting a significant improvement over the external classifier performance can be achieved using relational information.
- For parameter selection, the entropy based selection scheme was observed to be more robust compared to the maximum probability selection scheme.
- With respect to choice of C using inaccurate CV estimates, the DIR, LSR and WvRN-V2 methods were observed to be more robust.
- Overall, the probabilistic methods fared better compared to the function estimation methods. Within the probabilistic methods, the IR methods were competitive to the other methods. Within the proposed methods, the LSR, WvRN-V2 and LGC ranked better, in that order.
- In terms of speed, the original IR method was quite slow compared to other methods. Although the

DIR method was faster, it was still not competitive in speed to the proposed methods. Among the proposed methods, the WvRN-extensions, LGC and LSR methods were faster, ranked by speed in that order.

- Overall, considering the issues of speed, robustness and accuracy together, the LSR and WvRN-V2 methods performed the best.

References

- [1] Y. Bengio, O. Delalleau, and N. L. Roux. Label propagation and quadratic criterion. In O. Chapelle, B. Schölkopf, and A. Zien, editors, *Label Propagation and Quadratic Criterion*. MIT Press, 2006.
- [2] A. Corduneanu and T. Jaakkola. Distributed information regularization on graphs. In *NIPS*, pages 297–304, 2005.
- [3] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relations: Metric labeling and markov random fields. In *FOCS*, pages 14–23, 1999.
- [4] S. A. Mackassy and F. Provost. Classification in networked data: A toolkit and a univariate case study. *JMLR*, 8:935–983, 2007.
- [5] P. Sen, G. M. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective classification in network data. Technical Report CS-TR-4905, University of Maryland, 2008.
- [6] K. Tsuda. Propagating distributions on a hypergraph by dual information regularization. In *ICML*, 2005.
- [7] J. Wang, T. Jebara, and S.-F. Chang. Graph transduction via alternating minimization. In *ICML*, 2008.
- [8] D. Zhou, O. Bousquet, T. Lal, J. Weston, and B. Scholkopf. Learning with local and global consistency. In *NIPS*, pages 321–328, 2004.
- [9] X. Zhu, Z. Ghahramani, and J. Lafferty. Semi-supervised learning using gaussian fields and harmonic function. In *ICML*, pages 912–919, 2003.

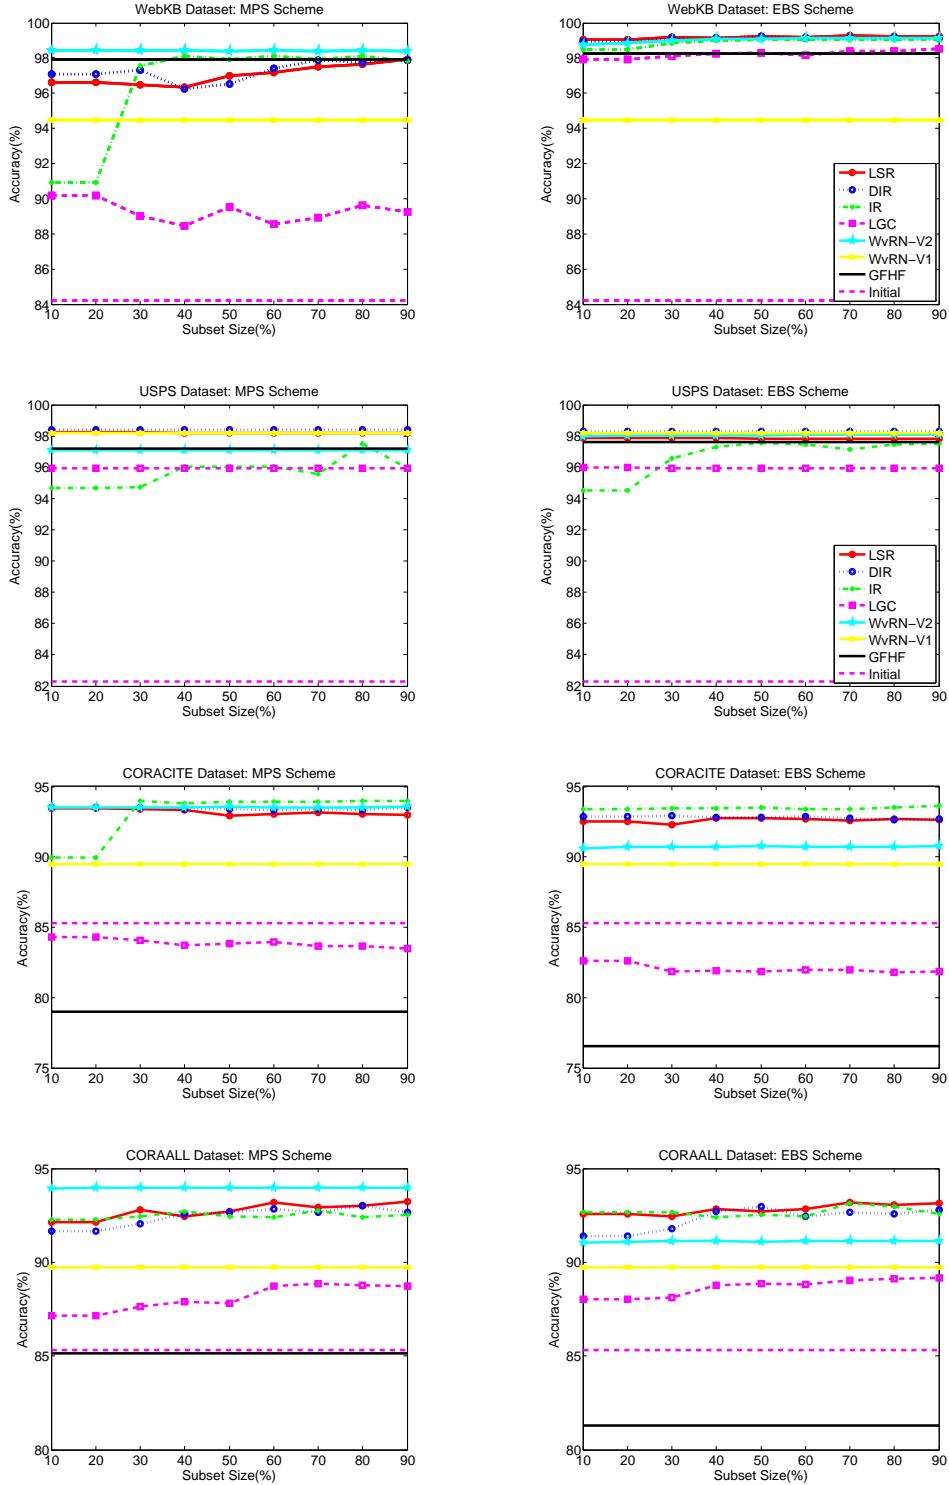


Figure 2: Setting 2 Results: X-axis represents the subset size in terms of percentage of number of nodes and Y-axis represents accuracy. The left and right columns correspond to MPS and EBS subset selection schemes respectively. The legends for the CORA datasets are same as the one given for the WebKB and USPS datasets. *Initial* refers to accuracy computed from generated inaccurate distribution.