

On the Effectiveness of Laplacian Normalization for Graph Semi-supervised Learning

Rie Johnson

*IBM T.J. Watson Research Center
Hawthorne, NY 10532, USA*

RIE1@US.IBM.COM

Tong Zhang

*Yahoo! Inc.
New York City, NY 10011, USA*

TZHANG@YAHOO-INC.COM

Editor: Charles Elkan

Abstract

This paper investigates the effect of Laplacian normalization in graph-based semi-supervised learning. To this end, we consider multi-class transductive learning on graphs with Laplacian regularization. Generalization bounds are derived using geometric properties of the graph. Specifically, by introducing a definition of graph cut from learning theory, we obtain generalization bounds that depend on the Laplacian regularizer. We then use this analysis to better understand the role of graph Laplacian matrix normalization. Under assumptions that the cut is small, we derive near-optimal normalization factors by approximately minimizing the generalization bounds. The analysis reveals the limitations of the standard degree-based normalization method in that the resulting normalization factors can vary significantly within each connected component with the same class label, which may cause inferior generalization performance. Our theory also suggests a remedy that does not suffer from this problem. Experiments confirm the superiority of the normalization scheme motivated by learning theory on artificial and real-world data sets.

Keywords: transductive learning, graph learning, Laplacian regularization, normalization of graph Laplacian

1. Introduction

Graph-based methods, such as spectral embedding, spectral clustering, and semi-supervised learning, have drawn much attention in the machine learning community. While various ideas have been proposed based on different intuitions, only recently have there been theoretical studies trying to understand why these methods work.

In spectral clustering, a traditional starting point is to find a partition of a graph that minimizes a certain definition of “graph cut” that quantifies the quality of the partition. The cut is the objective one attempts to minimize. Spectral methods can then be derived as a certain continuous relaxation that approximately solves the “graph cut” problem. Based on various intuitions and heuristics, various definitions of cuts have been proposed in the literature (for example, Shi and Malik, 2000, Ding et al., 2001, among others). In order to understand such methods, we need to ask the following two questions. First what is the quality of the relaxation approach as an approximation method to solve the original “graph cut” problem. Second, and more importantly, we need to understand why one should optimize one definition of “cut” instead of other alternatives. In the literature, different arguments and intuitions have been proposed to justify different choices. However, without a more

universally acceptable criterion, it is difficult to argue that one cut definition is better than another just based on heuristics. If a universally agreeable standard does exist, then one should focus on that criterion instead of an artificially defined cut problem.

For example, in the context of spectral clustering, there are two well-known types of graph cut, the ratio cut (Hagen and Kahng, 1992) and the normalized cut (Shi and Malik, 2000). Approximate optimization of the ratio cut leads to eigenvector computation of the *unnormalized* graph Laplacian matrix (which we will define later), and that of the normalized cut involves the *normalized* graph Laplacian matrix (normalized using node degrees). Although a number of empirical studies indicate that the normalized cut often leads to better clustering results, there isn't any direct theoretical proof except for some implicit evidence. As another example, the definition of graph Laplacian in the spectral graph theory in Chung (1998) is normalized, but that is for graph theoretical reasons instead of statistical reasons. Specifically, the normalized Laplacian allows easier translation of results from differential geometry, and it also allows consistent relations with conductance on Markov chains. The compatibility of continuous Laplacian on manifold and normalized graph Laplacian was also noted by von Luxburg et al. (2005) from a different perspective. Similarly, some analysis of spectral clustering employs normalized cut (Meilă et al., 2005, for example) as that makes derivation easier. These can be regarded as implicit evidence for preferring the normalized Laplacian over the unnormalized Laplacian. However, it has not been directly proved that such degree-based normalization (corresponding to normalized cut) should improve performance.

In order to understand this issue better, we take a different approach in this paper. Observe that for spectral clustering applications, there are often pre-defined (but unknown) clusters (classes). In this setting, the goal is to find such classes either using unsupervised or semi-supervised methods. Therefore for such problems, a universally agreeable standard is to find clusters that overlap significantly with the underlying class labels. That is, instead of using any artificially defined cut, we should design an algorithm to minimize the classification error. This is the criterion we focus on in this paper. We will see that normalization comes naturally into the generalization analysis we develop. By optimizing the corresponding generalization bounds, we seek to obtain a better understanding of the effect of Laplacian normalization.

In spectral clustering or graph based semi-supervised learning, one starts with *similarity graphs* that link similar data points. For example, one may connect data points that are close in the feature space to form a k -nearest neighbor graph. If the graph is fully connected within each class and disconnected between the classes, then appropriate cut minimization leads to perfect classification. It was proposed in Ng et al. (2002) that one may first project these data points into the eigenspace corresponding to the largest eigenvalues of a normalized adjacency matrix of the graph and then use the standard k -means method to perform clustering. The basic motivation is quite similar to that of Shi and Malik (2000). It can be shown that in the ideal case (each class forms a connected subgraph, and there is no inter-class edge), points in the same cluster will be mapped into a single point in the reduced eigenspace, while points in different clusters will be mapped to different points. This implies that for clustering the distance in the reduced space is better than the original distance. A natural question in this setting is how to design a distance function that leads to better clustering. While the argument in Ng et al. (2002) gives a satisfactory answer in the idealized case, it is far less clear what happens in general. One approach to address this problem is to learn a distance metric that can lead to more desirable clustering results from a set of labeled examples (for example, as in Xing et al., 2003). The inner product associated with a distance metric can be viewed as a kernel,

and the kernel fully determines the outcome of the k -means algorithm. Therefore this approach can also be viewed as designing a kernel optimal for clustering.

Closely related to clustering, one may also consider kernel design methods in semi-supervised learning using a discriminative method such as SVM (e.g., Lanckriet et al., 2004). In this setting, the change of the distance metric becomes a change of the underlying kernel. If the kernel is induced from a graph, then one may formulate semi-supervised learning directly on the graph; for example, see Belkin and Niyogi (2004), Szummer and Jaakkola (2002), Zhou et al. (2004) and Zhu et al. (2003). In these studies, the kernel is induced from the adjacency matrix \mathbf{W} whose (i, j) -entry is the weight of edge (i, j) . \mathbf{W} is often normalized by $\mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$ as in Chung (1998), Shi and Malik (2000), Ng et al. (2002), Zhou et al. (2004), where \mathbf{D} is a diagonal matrix whose (j, j) -entry is the degree of the j -th node, but sometimes not (Belkin and Niyogi, 2004, Zhu et al., 2003). Although such normalization can significantly affect the performance, the issue has not been carefully studied from the learning theory perspective. The relationship of kernel design and graph learning was investigated in Zhang and Ando (2006), where it was argued that quadratic regularization-based graph learning can be regarded as kernel design in the spectral domain. That is, one keeps the kernel eigenvectors and modifies the corresponding eigenvalues. Moreover if input data are corrupted with noise, then such spectral kernel design can help to improve classification performance. However, the analysis does not explain why normalization of the adjacency matrix \mathbf{W} is useful for practical purposes.

Our goals here are twofold. First we present a model for transductive learning on graphs and develop a margin analysis for multi-class graph learning. We then analyze graph learning using graph properties such as graph-cut and a concept we call *pure subgraph*. The analysis naturally employs quantities formalizing the standard graph-learning assumption that well connected nodes are likely to have the same label. Second, we use this analysis to obtain a better understanding of normalizing the Laplacian matrix $(\mathbf{D} - \mathbf{W})$ in graph semi-supervised learning. As mentioned above, normalization has been commonly practiced and appears to be useful, but there hasn't been any solid theoretical justification on why it should be useful. Our analysis addresses this issue from a learning theoretical point of view, and reveals a limitation of the standard degree-based normalization scheme. We then propose a remedy based on the learning theory results and use experiments to demonstrate that the remedy leads to improved classification performance.

This paper expands on our preliminary results reported in Ando and Zhang (2007).

2. Transductive Learning Model

We consider the following multi-category transductive learning model defined on a graph. Let $V = \{v_1, \dots, v_m\}$ be a set of m nodes, and let \mathcal{Y} be a set of K possible output values. Assume that each node v_j is associated with an output value $y_j \in \mathcal{Y}$, which we are interested in predicting. In order to do so, we randomly draw a set of n indices $Z_n = \{j_i : 1 \leq i \leq n\}$ from $\{1, \dots, m\}$ uniformly and without replacement. We manually label the n nodes v_{j_i} with labels $y_{j_i} \in \mathcal{Y}$, and then automatically label the remaining $m - n$ nodes. The goal is to estimate the labels on the remaining $m - n$ nodes as accurately as possible.

In this paper, we shall assume that the labels $y = [y_1, \dots, y_m]$ are deterministic. However, the analysis can also be applied if we have random labels. In the transductive learning setting considered in this paper, we may assume that we are given a single random draw $y = [y_1, \dots, y_m]$, which we fix. With this fixed y vector, we are interested in the performance of reconstructing it from a

subset of labels. This formulation is more appropriate for problems such as classification on graphs considered here.

In modern machine learning, instead of estimating the labels y_j directly, y_j is often encoded into a vector in R^K , so that the problem becomes that of generating an estimation vector $\mathbf{f}_j = [f_{j,1}, \dots, f_{j,K}] \in R^K$, which can then be used to recover the label y_j . In multi-category classification with K classes $\mathcal{Y} = \{1, \dots, K\}$, we encode each $y_j = k \in \mathcal{Y}$ as $e_k \in R^K$, where e_k is a vector of zero entries except for the k -th entry being one. Given a function $\mathbf{f}_j = [f_{j,1}, \dots, f_{j,K}] \in R^K$ (which is intended to approximate e_{y_j}), we decode the corresponding label estimation \hat{y}_j as:

$$\hat{y}_j = \hat{y}(\mathbf{f}_j) = \arg \max_k \{f_{j,k} : k = 1, \dots, K\}.$$

If the true label is y_j , then the corresponding classification error is:

$$\mathbf{err}(\mathbf{f}_j, y_j) = I(\hat{y}(\mathbf{f}_j) \neq y_j),$$

where we use $I(\cdot)$ to denote the set indicator function.

In order to estimate the concatenated vector $\mathbf{f} = [\mathbf{f}_j] = [f_{j,k}] \in R^{mK}$ from only a subset of labeled nodes, we have to impose restrictions on possible values of \mathbf{f} . In this paper, we consider restrictions defined through a quadratic regularizer of the following form:

$$\mathbf{f}^T \mathbf{Q}_K \mathbf{f} = \sum_{k=1}^K \mathbf{f}_{\cdot,k}^T \mathbf{K}^{-1} \mathbf{f}_{\cdot,k},$$

where $\mathbf{K} \in R^{m \times m}$ is a positive definite kernel matrix and $\mathbf{f}_{\cdot,k} = [f_{1,k}, \dots, f_{m,k}] \in R^m$. That is, the predictive vector for each class k is regularized separately. We assume that the kernel matrix \mathbf{K} is full-rank. We will consider the kernel matrix induced by the graph Laplacian, which we shall define later in the paper. Note that we use the bold symbol \mathbf{K} to denote the kernel matrix and the regular capitalized K to denote the number of classes.

Given a vector $\mathbf{f} \in R^{mK}$, the accuracy of its component $\mathbf{f}_j = [f_{j,1}, \dots, f_{j,K}] \in R^K$ is measured by a loss function $\phi(\mathbf{f}_j, y_j)$. Our learning method attempts to minimize the empirical risk on the set Z_n of n labeled training nodes, subject to $\mathbf{f}^T \mathbf{Q}_K \mathbf{f}$ being small:

$$\hat{\mathbf{f}}(Z_n) = \arg \min_{\mathbf{f} \in R^{mK}} \left[\frac{1}{n} \sum_{j \in Z_n} \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q}_K \mathbf{f} \right]. \tag{1}$$

where $\lambda > 0$ is an appropriately chosen regularization parameter.

In this paper, we focus on a special class of loss functions of the form $\phi(\mathbf{f}_j, y_j) = \sum_{k=1}^K \phi_0(f_{j,k}, \delta_{k,y_j})$, where $\delta_{a,b}$ is the delta function defined as: $\delta_{a,b} = 1$ when $a = b$ and $\delta_{a,b} = 0$ otherwise. In addition, we introduce the following assumption for convenience.

Assumption 1 Let $\phi(\mathbf{f}_j, y_j) = \sum_{k=1}^K \phi_0(f_{j,k}, \delta_{k,y_j})$ in (1), where $\mathbf{f}_j = [f_{j,1}, \dots, f_{j,K}] \in R^K$. Assume that there exist positive constants a , b , and c such that

- $\phi_0(x, y)$ is non-negative and convex in x .
- When $y = 0, 1$, and $\phi_0(x, y) \leq a$, $|\nabla_1 \phi_0(x, y)| \leq b$, where $\nabla_1 \phi_0(x, y)$ denotes a sub-gradient of $\phi_0(x, y)$ with respect to x .

- $c = \inf\{x : \phi_0(x, 1) \leq a\} - \sup\{x : \phi_0(x, 0) \leq a\}$.

The formulation presented here corresponds to the one-versus-all method for multi-category classification, and standard binary loss functions such as least squares, logistic regression, and SVMs can be used. For the SVM loss function $\phi_0(x, y) = \max(0, 1 - (2x - 1)(2y - 1))$, we may take $a = 0.5$, $b = 2$, and $c = 0.5$. For the least squares function $\phi_0(x, y) = (x - y)^2$, we may choose $a = 1/16$, $b = 0.5$, $c = 0.5$.

We are interested in the generalization behavior of (1) compared to a properly defined optimal regularized risk. This type of inequality is often referred to as ‘‘oracle inequality’’ in the learning theory literature and is particularly useful for analyzing the quality of the underlying learning method. The following theorem gives an oracle inequality, and its proof can be found in Appendix A.

Theorem 1 Consider (1) with loss function ϕ satisfying Assumption 1. Then $\forall p > 0$, the expected generalization error of the learning method (1) over the training samples Z_n , uniformly drawn without replacement from graph nodes $\{1, \dots, m\}$, can be bounded by:

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \text{err}(\hat{\mathbf{f}}_j(Z_n), y_j) \leq \frac{1}{a} \inf_{\mathbf{f} \in R^{m \times K}} \left[\frac{1}{m} \sum_{j=1}^m \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q} \mathbf{K} \mathbf{f} \right] + \left(\frac{b \text{tr}_p(\mathbf{K})}{\lambda n c} \right)^p,$$

where $\bar{Z}_n = \{1, \dots, m\} - Z_n$,

$$\text{tr}_p(\mathbf{K}) = \left(\frac{1}{m} \sum_{j=1}^m \mathbf{K}_{j,j}^p \right)^{1/p},$$

and $\mathbf{K}_{j,j}$ denotes the j -th diagonal entry of matrix \mathbf{K} .

If we take $p = 1$ in Theorem 1, then the bound depends on the trace of matrix \mathbf{K} : $\text{tr}(\mathbf{K}) = m \text{tr}_1(\mathbf{K})$. The trace of a kernel matrix has been employed in a number of previous studies to characterize generalization ability of kernel methods. The generalized quantity in Theorem 1 with $p \neq 1$ has non-trivial consequences which we will investigate in the paper.

Although we consider a specific form of loss function in this paper, one can obtain similar bounds with other forms of loss functions such as $\phi(\mathbf{f}_j, y_j) = \sup_{k \neq y_j} \phi_0(f_{j,y_j} - f_{j,k})$. What is important in our analysis are the two quantities $\mathbf{f}^T \mathbf{Q} \mathbf{K} \mathbf{f}$ and $\text{tr}_p(\mathbf{K})$ that determine the generalization performance. We will focus on the interpretation of these quantities.

3. Margin and Graph Cut

Consider an undirected graph $G = (V, E)$ defined on the nodes $V = \{v_j : j = 1, \dots, m\}$, with edges $E \subset \{1, \dots, m\} \times \{1, \dots, m\}$, and weights $w_{j,j'} \geq 0$ associated with edges $(j, j') \in E$. For simplicity, we assume that $(j, j) \notin E$ and $w_{j,j'} = 0$ when $(j, j') \notin E$. Let $\text{deg}_j(G) = \sum_{j'=1}^m w_{j,j'}$ be the degree of node j of graph G . We consider the following definition of normalized Laplacian.

Definition 2 Consider a graph $G = (V, E)$ of m nodes with weights $w_{j,j'}$ ($j, j' = 1, \dots, m$). The unnormalized Laplacian matrix $\mathcal{L}(G) \in R^{m \times m}$ is defined as: $\mathcal{L}_{j,j'}(G) = -w_{j,j'}$ if $j \neq j'$; $\text{deg}_j(G)$ otherwise. Given m scaling factors \mathbf{S}_j ($j = 1, \dots, m$), let $\mathbf{S} = \text{diag}(\{\mathbf{S}_j\})$. The \mathbf{S} -normalized Laplacian matrix is defined as: $\mathcal{L}_{\mathbf{S}}(G) = \mathbf{S}^{-1/2} \mathcal{L}(G) \mathbf{S}^{-1/2}$. The corresponding regularization is based on:

$$\mathbf{f}_{\cdot,k}^T \mathcal{L}_{\mathbf{S}}(G) \mathbf{f}_{\cdot,k} = \frac{1}{2} \sum_{j,j'=1}^m w_{j,j'} \left(\frac{f_{j,k}}{\sqrt{\mathbf{S}_j}} - \frac{f_{j',k}}{\sqrt{\mathbf{S}_{j'}}} \right)^2.$$

A common choice of \mathbf{S} is $\mathbf{S} = \mathbf{I}$, corresponding to regularizing with the unnormalized Laplacian \mathcal{L} . The idea is natural: we assume that the predictive values $f_{j,k}$ and $f_{j',k}$ should be close when $(j, j') \in E$ with a strong link. Another common choice is to normalize by $\mathbf{S}_j = \deg_j(G)$, as in Ng et al. (2002), Shi and Malik (2000), Zhou et al. (2004) and Chung (1998), which we refer to as degree-based normalization. At first sight, the need for normalization is not immediately clear. However, as we will show later, normalization using appropriate scaling factors can improve performance.

3.1 Generalization Analysis Using Graph-Cut

We will adapt Theorem 1 in Section 2 to analyze graph learning using graph properties such as graph-cut. We now introduce a learning theoretical definition of \mathbf{S} -normalized graph cut as follows.

Definition 3 Given label $y = \{y_j\}_{j=1,\dots,m}$ on V , we define the cut for the \mathbf{S} -normalized Laplacian $\mathcal{L}_{\mathbf{S}}$ in Definition 2 as:

$$\text{cut}(\mathcal{L}_{\mathbf{S}}, y) = \sum_{j, j': y_j \neq y_{j'}} \frac{w_{j, j'}}{2} \left(\frac{1}{\mathbf{S}_j} + \frac{1}{\mathbf{S}_{j'}} \right) + \sum_{j, j': y_j = y_{j'}} \frac{w_{j, j'}}{2} \left(\frac{1}{\sqrt{\mathbf{S}_j}} - \frac{1}{\sqrt{\mathbf{S}_{j'}}} \right)^2.$$

Note that unlike typical graph-theoretical definitions of graph-cut in the literature, the learning theoretical definition of cut not only penalizes a normalized version of between-class edge weights, but also penalizes within-class edge weights when such an edge connects two nodes with different scaling factors. This difference has important consequences, which we will investigate later in the paper. For unnormalized Laplacian, the second term on the right hand side of Definition 3 vanishes, which means that it only penalizes weights corresponding to edges connecting nodes with different labels. In this case, the learning theoretical definition corresponds to the graph-theoretical definition: $\text{cut}(\mathcal{L}, y) = \sum_{j, j': y_j \neq y_{j'}} w_{j, j'}$.

It is worth noting that in our framework, cut is used to indicate the absolute amount of perturbation from the idealized case with zero-cut. In spectral clustering, the absolute cut is often scaled and the resulting quantity is used as a quality measure for the clusters. In comparison, our quality measure is always to minimize the classification error. In particular, the unnormalized Laplacian is used in spectral clustering to approximately minimize the *ratio cut* $= \sum_{j \in A, j' \in B} w_{j, j'} / (|A| \cdot |B|)$ (Hagen and Kahng, 1992) instead of $\sum_{j \in A, j' \in B} w_{j, j'}$. The scaling in the ratio cut (when $K = 2$) corresponds to the normalization of a specific encoding of the target vectors ($\mathbf{f}_{\cdot, k} \in \mathbb{R}^m$ which encodes the true output values in our setting); it is used in the computation of the second smallest eigenvalue of the unnormalized Laplacian. Our analysis throughout this paper assumes unmodified target vectors with components taking values in $\{0, 1\}$, which differs from Hagen and Kahng (1992). As such, the discrepancy of cut definition is merely due to a difference in representation of the target vectors, and both lead to the same unnormalized Laplacian matrix algorithmically. Although our definition is motivated by classification problems, it may also be useful for clustering problems because there is a strong relationship between clustering and transductive classification (where class labels indicate the underlying clusters). This means that the normalization method we propose in the paper might be useful in spectral clustering as well.

Using the learning theoretical graph-cut definition, we can obtain a generalization result for the estimator in (1) with \mathbf{K} defined as follows:

$$\mathbf{K}^{-1} = \alpha \mathbf{S}^{-1} + \mathcal{L}_{\mathbf{S}}(G) = \mathbf{S}^{-1/2}(\alpha \mathbf{I} + \mathcal{L}(G))\mathbf{S}^{-1/2}, \quad (2)$$

where \mathbf{I} is the identity matrix. Note that $\alpha > 0$ is a tuning parameter to ensure that \mathbf{K} is strictly positive definite. As we will see later, this parameter is important. The corresponding regularization condition is

$$\mathbf{f}^T \mathbf{Q}_{\mathbf{K}} \mathbf{f} = \sum_{k=1}^K \left(\alpha \sum_{j=1}^m \frac{f_{k,j}^2}{\mathbf{S}_j} + \frac{1}{2} \sum_{j,j'=1}^m \left(\frac{f_{j,k}}{\sqrt{\mathbf{S}_j}} - \frac{f_{j',k}}{\sqrt{\mathbf{S}_{j'}}} \right)^2 w_{j,j'} \right).$$

Another possibility is to let $\mathbf{K}^{-1} = \alpha \mathbf{I} + \mathcal{L}_{\mathbf{S}}(G)$. The conclusions, which we will not include in this paper, are similar to that of (2).

For simplicity, we state the generalization bound based on Theorem 1 with optimal λ . Note that in applications, $\hat{\lambda}$ is usually tuned through cross validation. Therefore assuming optimal λ will simplify the bound so that we can focus on the more essential characteristics of generalization performance. The following assumption is used to simplify the bound

Assumption 2 Consider (1) with regularization condition (2), loss function ϕ satisfying Assumption 1, and assume that $\phi_0(0,0) = \phi_0(1,1) = 0$.

It is easy to check that the conditions on the loss function in Assumption 2 hold for the least squares method (which we focus on in this paper) as well as other standard loss functions such as SVM.

Theorem 4 Consider (1) such that Assumption 2 is satisfied. Then $\forall p > 0$, there exists a sample independent regularization parameter $\hat{\lambda}$ in (1) such that the expected generalization error is bounded by:

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j(Z_n), y_j) \leq \frac{C_p(a,b,c)}{n^{p/(p+1)}} (\alpha s + \mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y))^{p/(p+1)} \mathbf{tr}_p(\mathbf{K})^{p/(p+1)},$$

$$C_p(a,b,c) = (b/ac)^{p/(p+1)} (p^{1/(p+1)} + p^{-p/(p+1)}), \quad (3)$$

where $s = \sum_{j=1}^m \mathbf{S}_j^{-1}$.

Proof Let $f_{j,k} = \delta_{y_j,k}$. It can be easily verified that

$$\frac{1}{m} \sum_{j=1}^m \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q}_{\mathbf{K}} \mathbf{f} = \lambda (\alpha s + \mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)).$$

Now, using this expression in Theorem 1, and then optimizing over λ , we obtain the desired inequality. \blacksquare

Note that with the least squares loss, we can take $b/ac = 16$ in Theorem 1. With a fixed p , the generalization error decreases at the rate $O(n^{-p/(p+1)})$ when the sample size n increases. This rate of convergence is faster when p increases. However in general, $\mathbf{tr}_p(\mathbf{K})$ is an increasing function of p . Therefore we have a trade-off between the two terms, and without appropriate normalization (which

we will consider later in the paper), one may prefer a smaller p in order to optimize the bound. An analysis will be provided in the next section. The bound also suggests that if we normalize \mathbf{K} so that its diagonal entries $\mathbf{K}_{j,j}$ become a constant, then $\text{tr}_p(\mathbf{K})$ is independent of p , and thus a larger p can be used in the bound. This motivates the idea of normalizing the diagonals of \mathbf{K} , which we will further investigate later in the paper. The generalization bound in Theorem 4 is closely related to the margin analysis for binary linear classification. Specifically, the right hand side can be viewed as a margin-like-quantity associated with the target function $f_{j,k} = \delta_{y_j,k}$ that separates the data. Here it is related to the concept of graph cut. Our goal is to better understand the quantity $(\alpha s + \text{cut}(\mathcal{L}_S, y))^{p/(p+1)} \text{tr}_p(\mathbf{K})^{p/(p+1)}$ using graph properties, which gives better understanding of graph based learning.

In the following, we will give example applications of Theorem 4. They illustrate that theoretically it is important to tune the parameter α to achieve good performance, which is also empirically observed in our experiments.

3.2 Zero-cut and Geometric Margin Separation

We consider an application of Theorem 4 for the unnormalized Laplacian under the zero-cut assumption that each connected component of the graph has a single label. With this assumption, the task is simply to estimate what label each connected component has.

Theorem 5 *Consider (1) such that Assumption 2 is satisfied and the regularization condition is $\mathbf{K}^{-1} = \alpha \mathbf{I} + \mathcal{L}$. Assume that $\text{cut}(\mathcal{L}, y) = 0$, and the graph has q connected components of sizes $m_1 \leq \dots \leq m_q$ ($\sum_{\ell} m_{\ell} = m$). For all $p > 0$, let $\alpha \rightarrow 0$, and with optimal λ , we have the generalization bound*

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in Z_n} \text{err}(\hat{\mathbf{f}}_j, y_j) \leq \frac{C_p(a, b, c)}{n^{p/(p+1)}} \left(\sum_{\ell=1}^q (m/m_{\ell})^{p-1} \right)^{1/(p+1)} + O(\alpha),$$

where C_p is defined in (3). In particular, we have

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in Z_n} \text{err}(\hat{\mathbf{f}}_j, y_j) \leq \min \left[2\sqrt{\frac{b}{ac} \cdot \frac{q}{n}}, \frac{b}{ac} \cdot \frac{m}{nm_1} \right] + O(\alpha).$$

Proof Since the graph has q connected components, \mathcal{L} has q eigenvectors \mathbf{v}_{ℓ} ($\ell = 1, \dots, q$) associated with zero-eigenvalues, where each eigenvector \mathbf{v}_{ℓ} is the indicator function of the ℓ -th connected component in the graph, that is, the j -th entry of vector \mathbf{v}_{ℓ} is 1 if j belongs to the ℓ -th connected component and 0 otherwise. It is not hard to check that as $\alpha \rightarrow 0$, $\alpha \mathbf{K} \rightarrow \sum_{\ell=1}^q \frac{1}{m_{\ell}} \mathbf{v}_{\ell} \mathbf{v}_{\ell}^T + O(\alpha)$. Therefore $\alpha \text{tr}_p(\mathbf{K}) \rightarrow m^{-1/p} (\sum_{\ell=1}^q m_{\ell}^{1-p})^{1/p}$. Now, we can use Theorem 4 to obtain the first inequality. The second inequality is obtained by setting $p = 1$ and by letting $p \rightarrow \infty$ on the right hand side. ■

Under the zero-cut assumption, the generalization performance can be bounded as $O(\sqrt{q/n})$ when $\alpha \rightarrow 0$. However, we can also achieve a faster convergence rate of $O(1/n)$, although the generalization performance depends on the inverse of the smallest component size through $m/m_1 \geq$

q . This implies that we will achieve better convergence at the $O(1/n)$ level if the sizes of the components are balanced. If the component sizes are significantly different, the convergence may behave like $O(\sqrt{q/n})$.

We discuss a concrete example in which Theorem 5 is applicable. Assume that each node v_j is associated with a data point x_j that belongs to the d -dimensional unit ball $B = \{x \in R^d : \|x\|_2 \leq 1\}$. We form a graph by connecting all nodes v_j to their nearest neighbors. In particular, we may consider an ϵ -ball centered at each v_j : $B_j(\epsilon) = \{x : \|x - x_j\|_2 \leq \epsilon\}$. We then form a graph by connecting each j with all points within the ball $B_j(\epsilon)$ and with unit weights.

We say that the data points are separable with geometric margin γ if for each node v_j the ball $B_j(\gamma)$ only contains points in class y_j . Now assume we use a ball of size $\epsilon \leq \gamma$. In this case, $\text{cut}(\mathcal{L}, y) = 0$, and there is a constant $q \leq \epsilon^{-d}$ such that the graph has at most q connected components, and we have:

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \text{err}(\hat{\mathbf{f}}_j, y_j) \leq 2\sqrt{\frac{b}{ac} \cdot \frac{q}{n}} + O(\alpha).$$

This bound does not depend on margin γ but depends only on q , the number of connected components. So even if the margin γ is small, the bound can still be good as long as q is small. This result can be used to understand why graph based semi-supervised learning may work better than standard kernel learning. In fact, it is not possible to derive similar generalization bounds for supervised learning because one needs unlabeled data (in addition to labeled data) to define such connected components. This means that graph semi-supervised learning can take advantage of the new quantity q to characterize its generalization performance, and this quantity cannot be used by standard supervised learning.

Note that we have assumed a very specific generative model for the data. In particular, if the data are generated in a way such that the number of connected components q is small, and each connected component belongs to a single class, then graph based semi-supervised learning can work better than supervised kernel learning. If this assumption does not hold (at least approximately), then graph based learning methods may fail. However, for many practical applications, the geometric margin separation assumption does appear quite reasonable. Therefore for such problems, graph based semi-supervised learning, which can take advantage of the underlying data generation model, may become helpful.

This section only considers a special case where the graph has q connected components. In this particular situation, the learning method (1) and the analysis provided here may not be optimal. The best method is just to identify each connected component to be a cluster and then determine its label by looking at one point of the cluster. However, this idea won't generalize to graphs with components that are weakly connected. In comparison, our analysis can easily generalize to that situation, as we shall investigate in the next section.

3.3 Non-Zero Cut and Pure Components

It is often too restrictive to assume that each connected component has only one label (that is, the cut is zero). In this section, we show that similar bounds can be obtained when this data generation assumption is relaxed. We are still interested in giving a characterization of the performance of (1) in terms of properties of the graph and introduce the following definition.

Definition 6 A subgraph $G_0 = (V_0, E_0)$ of $G = (V, E)$ is called a pure component if G_0 is connected, E_0 is induced by restricting E on V_0 , and the labels y have identical values on V_0 . A pure subgraph

$G' = \cup_{\ell=1}^q G_\ell$ of G divides V into q disjoint sets $V = \cup_{\ell=1}^q V_\ell$ such that each subgraph $G_\ell = (V_\ell, E_\ell)$ is a pure component. Denote by $\hat{\lambda}_i(G_\ell) = \lambda_i(\mathcal{L}(G_\ell))$ the i -th smallest eigenvalue of $\mathcal{L}(G_\ell)$.

For instance, if we remove all edges of G that connect nodes with different labels, then the resulting subgraph is a pure subgraph (though it may not be the only one). For each pure component G_ℓ , its first eigenvalue $\lambda_1(G_\ell)$ is always zero. The second eigenvalue $\lambda_2(G_\ell) > 0$ because G_ℓ is connected. This $\lambda_2(G_\ell)$ can be regarded as a measurement of how well G_ℓ is connected. We use it together with graph cut to derive a generalization bound. The proof is given in Appendix B.

Theorem 7 Consider (1) such that Assumption 2 is satisfied. Let $G' = \cup_{\ell=1}^q G_\ell$ ($G_\ell = (V_\ell, E_\ell)$) be a pure subgraph of G . For all $p \geq 1$, there exist sample-independent regularization parameter λ and a fixed tuning parameter α , such that

$$\begin{aligned} & \mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \\ & \leq \frac{C_p(a, b, c)}{n^{p/(p+1)}} \left(s^{1/2} \left(\sum_{\ell=1}^q \frac{s_\ell(p)/m}{m_\ell^p} \right)^{1/2p} + \mathbf{cut}(\mathcal{L}_S, y)^{1/2} \left(\sum_{\ell=1}^q \frac{s_\ell(p)/m}{\lambda_2(G_\ell)^p} \right)^{1/2p} \right)^{2p/(p+1)}, \end{aligned}$$

where C_p is defined in (3), $m_\ell = |V_\ell|$, $s = \sum_{j=1}^m \mathbf{S}_j^{-1}$, and $s_\ell(p) = \sum_{j \in V_\ell} \mathbf{S}_j^p$.

Theorem 7 is a natural generalization of Theorem 5 when $p \geq 1$. It quantitatively illustrates the importance of analyzing graph learning using a partition of the original graph into well-connected pure components. The second eigenvalue $\lambda_2(G_i)$ measures how well-connected G_i is. A more intuitive quantity that measures the connectedness of graph $G = (V, E)$ is the *isoperimetric number* h_G defined as

$$h_G = \inf_{S \subset V} \sum_{j \in S, j' \in V-S} w_{j,j'} / \min(|S|, |V-S|).$$

It is well-known that $\lambda_2(G_i) \geq h_{G_i}^2 / (2 \max_j \deg_j(G_i))$ (Chung, 1998). The isoperimetric number of a graph is large when the nodes are well-connected everywhere. In particular, if $\deg_j(G)$ is of the order $|V|$, and $w_{i,j} = 1$ when $(i, j) \in E$, then for a well-connected graph, $\sum_{j \in S, j' \in V-S} w_{j,j'}$ is of the order $|S||V-S|$, and $h_G = O(|V|)$. Let G' be a well-behaved pure-subgraph of G , such that each pure component G_ℓ of G' is well-connected in the above sense. We thus have the condition

$$\lambda_2(G_\ell)/m_\ell \geq u(G')$$

for some constant $u(G')$ that does not depend on the size of the pure components (but only how well-connected each pure component is). Under this condition, we may replace $\sum_{\ell=1}^q m_\ell \lambda_2(G_\ell)^{-p}$ by $u(G')^{-p} \sum_{\ell=1}^q m_\ell^{1-p}$ in Theorem 7 and obtain a simplified bound:

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq \frac{C_p(a, b, c)}{n^{p/(p+1)}} \left(\sum_{\ell=1}^q \frac{s_\ell(p)/m}{(m_\ell/m)^p} \right)^{1/(p+1)} \left(\sqrt{\frac{s}{m}} + \sqrt{\frac{\mathbf{cut}(\mathcal{L}_S, y)}{u(G')m}} \right)^{2p/(p+1)},$$

where we define $u(G') = \min_{\ell}(\lambda_2(G_{\ell})/m_{\ell})$. We consider two special cases: $p = 1$ and $p \rightarrow \infty$:

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq 2 \sqrt{\frac{b}{ac} \cdot \frac{\sum_{\ell=1}^q (s_{\ell}(1)/m_{\ell})}{n}} \left(\sqrt{\frac{s}{m}} + \sqrt{\frac{\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)}{u(G')m}} \right), \quad (4)$$

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq \frac{b}{ac} \cdot \frac{\max_{\ell} \max_{j \in V_{\ell}} (\mathbf{S}_j/m_{\ell})}{n} \left(\sqrt{s} + \sqrt{\frac{\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)}{u(G')}} \right)^2. \quad (5)$$

These bounds are generalizations of those in Theorem 5. Suppose that we take $\mathbf{S} = \mathbf{I}$. Then the number of pure components q affects the $O(1/\sqrt{n})$ convergence rate in (4) as $\sum_{\ell=1}^q s_{\ell}(1)/m_{\ell} = q$. If the sizes of the components are balanced, we can achieve better convergence at the $O(1/n)$ level as in (5); otherwise, the convergence may behave like $O(\sqrt{q/n})$. This observation motivates a scaling matrix \mathbf{S} that compensates for the unbalanced pure component sizes, which we will investigate next.

3.4 Optimal Normalization for Near-zero-cut Partition

As discussed in the introduction, the common practice of the normalization of the adjacency matrix (\mathbf{W}) or the graph Laplacian ($\mathbf{D} - \mathbf{W}$) is based on degrees, which corresponds to setting $\mathbf{S} = \mathbf{D}$. Although such normalization may significantly affect the performance, to our knowledge, there is no learning theory analysis on the effect of normalization. The purpose of this section is to fill this gap using the theoretical tools developed earlier. We shall focus on a near ideal situation to gain intuition.

Consider a pure subgraph $G' = \cup_{\ell=1}^q G_{\ell}$ ($G_{\ell} = (V_{\ell}, E_{\ell})$) of G . If scaling factors \mathbf{S}_j are approximately constant within each pure component, then using the Laplacian in Definition 2, we have a small regularization penalty for the edges within a pure component and between the nodes that have close output values (i.e., $f_{j,k} \approx f_{j',k}$). Therefore, in the following we focus on finding the optimal scaling matrix \mathbf{S} such that \mathbf{S}_j is constant within each pure component V_{ℓ} , and assume that \mathbf{S} is quantified by q numbers $[\bar{s}_{\ell}]_{\ell=1, \dots, q}$, such that $\mathbf{S}_j = \bar{s}_{\ell}$ when $j \in V_{\ell}$.

Consider the following quantity:

$$\mathbf{cut}(G', y) = \sum_{j, j': y_j \neq y_{j'}} w_{j, j'} + \sum_{\ell \neq \ell'} \sum_{j \in V_{\ell}, j' \in V_{\ell'}} \frac{w_{j, j'}}{2}.$$

It is easy to check that

$$\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y) \leq \mathbf{cut}(G', y) / \min_{\ell} \bar{s}_{\ell}.$$

Assume that weights are small between pure components, and therefore, $\mathbf{cut}(G', y)$ is small.

With the $O(1/n)$ convergence rate, we obtain from (5) that

$$\frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq \frac{b}{ac} \cdot \frac{\max_{\ell} (\bar{s}_{\ell}/m_{\ell})}{n} \left(\sqrt{\sum_{\ell=1}^q m_{\ell}/\bar{s}_{\ell}} + \sqrt{\frac{\mathbf{cut}(G', y)}{u(G') \min_{\ell} \bar{s}_{\ell}}} \right)^2.$$

If $\mathbf{cut}(G', y)$ is small, then the dominating term on the right hand side is

$$\frac{\max_{\ell} (\bar{s}_{\ell}/m_{\ell})}{n} \sum_{\ell=1}^q \frac{m_{\ell}}{\bar{s}_{\ell}},$$

which can be optimized with the choice $\bar{s}_\ell = m_\ell$, and the resulting bound becomes:

$$\frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq \frac{b}{ac} \cdot \frac{1}{n} \left(\sqrt{q} + \sqrt{\frac{\mathbf{cut}(G', y)}{u(G') \min_\ell m_\ell}} \right)^2.$$

That is, if $\mathbf{cut}(G', y)$ is small, then we can choose scaling factor $\bar{s}_\ell \propto m_\ell$ for each pure component ℓ so that the generalization performance is approximately $(ac)^{-1} b \cdot q/n$, which is of the order $O(1/n)$.

The analysis provided here not only proves the importance of normalization under the learning theoretical framework, but also suggests that the good normalization factor for each node j is approximately the size of the well-connected pure component that contains node j (assuming that nodes belonging to different pure components are only weakly connected). Our analysis focused on the case that the scaling factors are a constant within each pure component. This condition is quite natural if we look at the normalized Laplacian regularization condition in Definition 2, where $f_{j,k}/\sqrt{\mathbf{S}_j}$ should be similar to $f_{j',k}/\sqrt{\mathbf{S}_{j'}}$ when $w_{j,j'}$ is large. If j and j' belongs to the same class, then $f_{j,k}$ should be similar to $f_{j',k}$. Therefore for such a pair (j, j') , we want to have $\mathbf{S}_j \approx \mathbf{S}_{j'}$ if $w_{j,j'}$ is large. Note that this requirement is not enforced by the standard degree-based normalization method $\mathbf{S}_j = \deg_j(G)$ because a well-connected pure component may contain nodes with quite different degrees. The assumption is satisfied under a simplified ‘‘box model’’, which is related to the models used by some previous researchers to derive the standard normalization method (e.g., Shi and Malik, 2000). In this model, a pure component is completely connected, and each node connects to all other nodes and itself with edge weight $w_{j,j'} = 1$. The degree is thus $\deg_j(G_\ell) = |V_\ell| = m_\ell$, which gives the optimal scaling in our analysis.

In general, the box model may not be a good approximation for practical problems. A more realistic approximation, which we call core-satellite model, will be introduced in the experimental section. For such a model, the degree-based normalization can fail because the $\deg_j(G_\ell)$ within each pure component G_ℓ is not approximately constant, and it may not be proportional to m_ℓ . In general, this approximation using degrees causes \mathbf{S}_j to potentially vary significantly within a pure component because each \mathbf{S}_j is only determined by its immediate neighbors.

Our analysis suggests that it is necessary to modify the degree-based scaling method $\mathbf{S}_j = \deg_j(G)$ so that the scaling factor is approximately a constant within each pure component, which should be proportional to m_ℓ . Our remedy is to look for connected components at a larger distance scale. Although there could be various methods to achieve this effect, we shall focus on a specific method motivated by the proofs of Theorem 5 and Theorem 7. Let $\bar{\mathbf{K}} = (\alpha \mathbf{I} + \mathcal{L})^{-1}$ be the kernel matrix corresponding to the unnormalized Laplacian. Using the terminology in the proofs, we observe that for small α :

$$\alpha \bar{\mathbf{K}} = \sum_{\ell=1}^q \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell + O(1),$$

and thus $\bar{\mathbf{K}}_{j,j} \propto m_\ell^{-1}$ for each $j \in V_\ell$. Therefore with small α , the scaling factor $\mathbf{S}_j = 1/\bar{\mathbf{K}}_{j,j}$ is near optimal for all j . For $\alpha > 0$, the effect of this scaling factor is essentially equivalent to looking for connected components at a scale of at most $O(1/\alpha)$ number of nodes. We call this method of normalization *\mathbf{K} -scaling* in this paper. It is equivalent to a normalization of the kernel matrix \mathbf{K} , so that each $\mathbf{K}_{j,j} = 1$. Although this method coincides with a common practice in standard kernel learning, it is important to notice that to show this method behaves well in the graph learning setting is highly non-trivial and novel. To our best knowledge, no one has proposed this normalization

method in the graph learning setting before. In fact, without learning theoretical results developed in this paper, it is not obvious that this method should work better than the more standard degree-based normalization method. In our framework, the main advantage of \mathbf{K} -scaling (compared to the standard degree-scaling, which we call \mathbf{L} -scaling) is twofold:

- The resulting \mathbf{S}_j does not vary significantly within a well-connected pure component.
- The resulting scaling is approximately m_ℓ (at a scale of $1/\alpha$), which is predicted by our theory to be desirable.¹

The superiority of this method will be demonstrated in our experiments. The main drawback of this method is the computational cost of directly inverting $(\alpha\mathbf{I} + \mathcal{L})$. For large scale problems, approximation methods are required.

3.5 Dimension Reduction

Normalization and dimension reduction have been commonly used in spectral clustering such as Ng et al. (2002) and Shi and Malik (2000). For semi-supervised learning, dimension reduction (without normalization) is known to improve performance (Belkin and Niyogi, 2004, Zhang and Ando, 2006) while the degree-based normalization (without dimension reduction) has also been explored (Zhou et al., 2004). In this section, we present a brief high-level argument that an appropriate combination of normalization and dimension reduction (as commonly used in spectral clustering) can improve classification performance. Detailed analysis can be found in Appendix C.

Let us first introduce dimension reduction with normalized Laplacian $\mathcal{L}_\mathbf{S}(G)$. Denote by $\mathbf{P}_\mathbf{S}^r(G)$ the projection operator onto the eigenspace of $\alpha\mathbf{S}^{-1} + \mathcal{L}_\mathbf{S}(G)$ corresponding to the r smallest eigenvalues. Now, we may define the following regularizer on the reduced subspace:

$$\mathbf{f}_{\cdot,k}^T \mathbf{K}^{-1} \mathbf{f}_{\cdot,k} = \begin{cases} \mathbf{f}_{\cdot,k}^T \mathbf{K}_0^{-1} \mathbf{f}_{\cdot,k} & \text{if } \mathbf{P}_\mathbf{S}^r(G) \mathbf{f}_{\cdot,k} = \mathbf{f}_{\cdot,k}, \\ +\infty & \text{otherwise.} \end{cases} \quad (6)$$

The benefit of dimension reduction in graph learning has been investigated in Zhang and Ando (2006), under the spectral kernel design framework. The idea is to modify the kernel eigenvalues so that the target spectral coefficient matches the kernel coefficients. Note that the normalization issue, which will change the eigenvectors and their ordering, wasn't investigated there. However, with a fixed scaling matrix \mathbf{S} , the reasoning given in Zhang and Ando (2006) can also be applied here. It was shown there that if noise is added into the kernel matrix, then in general kernel eigenvalues will decay slower than the target spectral coefficients. Because of this, dimension reduction, which makes kernel eigenvalues better match the decay of target spectral coefficients, will become helpful. For Laplacian regularization investigated here, we may regard noise as edges connecting pure components of different classes, which increase the cut in Definition 3. Such noise can be significantly reduced if we project it into a low-dimensional space, and if the target functions approximately lie in this low-dimensional space. In this context, the effect of modification of eigenspaces through appropriate Laplacian normalization is to achieve faster decay of the target spectral coefficients in the

1. Although “the scaling factor $\mathbf{S}_j = m_\ell$ ” might be reminiscent of the ratio cut in spectral clustering, note that, as mentioned earlier, the ratio cut corresponds to the unnormalized Laplacian. \mathbf{K} -scaling suggested here normalizes the Laplacian matrix, which is in the same spirit of normalized cut (degree-scaling) but fixes some of its short-comings based on learning theoretical insights developed here.

	classes #1, #2	classes #3–#10
graph1	(4,2)	(2,1)
graph2	(6,3)	(2,1)
graph3	(8,4)	(2,1)

Figure 1: Generation of graph 1–5. (c, e) in the table indicates that for each node, we randomly chose c nodes of the same class and connect it to them, and we randomly chose e nodes of other classes (introducing errors) and connect it to them. Edge weights are fixed to 1.




Figure 2: Classification accuracy (%) on the graphs where degrees are nearly constant within the class. $n = 40, m = 2000$. With dimension reduction ($\dim \leq 20$; chosen by cross validation). Average over 10 random splits with one standard deviation.

first few eigenvectors of the kernel. Therefore, under certain conditions, dimension reduction can reduce noise (corresponding to a small cut), which essentially makes normalization more effective as shown in Section 3.4.

We show our formal analysis of the combination of dimension reduction (as in (6) above) and normalization of Laplacian, for completeness, in Appendix C and empirical results in the next section.

4. Experiments

We experiment with the Laplacian regularization with the normalization methods discussed above, on synthesized data sets generated by controlling graph properties as well as three real-world data sets.

4.1 Experimental Framework

The Laplacian matrix \mathcal{L} is generated from a graph G so that $\mathcal{L}_{j,j'} = -w_{j,j'}$ for $j \neq j'$ and $\mathcal{L}_{j,j} = \deg_j(G)$. Using \mathcal{L} , we define matrix \mathbf{K} as follows:

- *Unnormalized*: $\mathbf{K} = (\alpha \mathbf{I} + \mathcal{L})^{-1}$. That is, $\mathbf{S} = \mathbf{I}$. No scaling.

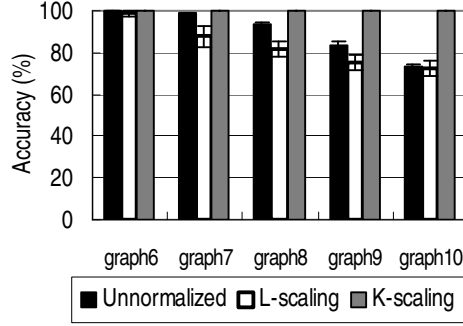


Figure 3: Classification accuracy on the core-satellite graphs. $n = 40, m = 2000$. With dimension reduction ($\dim \leq 20$; chosen by cross validation). Average over 10 random splits with one standard deviation.

- **K-scaling:** $\mathbf{K} = (\mathbf{S}^{-1/2}(\alpha\mathbf{I} + \mathcal{L})\mathbf{S}^{-1/2})^{-1}$ where $\mathbf{S} = \text{diag}_j(\bar{\mathbf{K}}_{j,j}^{-1})$ with $\bar{\mathbf{K}} = (\alpha\mathbf{I} + \mathcal{L})^{-1}$. The diagonal entries of \mathbf{K} are all ones.
- **L-scaling:** $\mathbf{K} = (\alpha\mathbf{I} + \mathbf{S}^{-1/2}\mathcal{L}\mathbf{S}^{-1/2})^{-1}$ where $\mathbf{S} = \text{diag}_j(\text{deg}_j(G))$. The diagonal entries of \mathbf{K}^{-1} are constant $(\alpha + 1)$. This is the standard degree-based scaling.

Using these three types of matrix \mathbf{K} , we test the following two types of regularization. One regularizes by $\mathbf{f}_{\cdot,k}^T \mathbf{K}^{-1} \mathbf{f}_{\cdot,k}$ using \mathbf{K} without dimension reduction, as in Section 3. The other reduces the dimension of \mathbf{K}^{-1} to r by leaving out all but several eigenvectors corresponding to the smallest r eigenvalues to obtain the eigenspace projector $\mathbf{P}_S^r(G)$ and regularizes by:

$$\begin{cases} \mathbf{f}_{\cdot,k}^T \mathbf{K}^{-1} \mathbf{f}_{\cdot,k} & \text{if } \mathbf{P}_S^r(G) \mathbf{f}_{\cdot,k} = \mathbf{f}_{\cdot,k} \\ +\infty & \text{otherwise} \end{cases}$$

as in Section 3.5. We use the one-versus-all strategy and use least squares as our loss function: $\phi_k(a, b) = (a - \delta_{k,b})^2$.

From m data points, n training labeled examples are randomly chosen while ensuring that at least one training example is chosen from each class. The remaining $m - n$ data points serve as test data. The regularization parameter λ is chosen by cross validation on the n training labeled examples. We will show performance when the rest of the parameters (α and dimensionality r) are also chosen by cross validation on the training labeled examples and when they are set to the optimum. The dimensionality r is chosen from $K, K + 5, K + 10, \dots, 100$ where K is the number of classes unless otherwise specified. Our focus is on small n close to the number of classes. Throughout this section, we conduct 10 runs with random training/test splits and report the average accuracy.

4.2 Controlled Data Experiments

The purpose of the controlled data experiments is to observe the correlation of the effectiveness of the normalization methods with graph properties. The graphs we generate contain 2000 nodes, each of which is assigned one of 10 classes.

First, we show the results when dimension reduction is applied to the three types of matrix \mathbf{K} . Figure 2 shows classification accuracy on three graphs that were generated so that the node degrees

(of either correct edges or erroneous edges) are close to constant within each class but vary across classes. Details of their generation are described in Figure 1. We observe that on these graphs, both \mathbf{K} -scaling and \mathbf{L} -scaling significantly improve classification accuracy over the unnormalized baseline. There is no prominent difference between \mathbf{K} -scaling’s and \mathbf{L} -scaling’s performance.

Observe that \mathbf{K} -scaling and \mathbf{L} -scaling perform differently on the graphs used in Figure 3. These graphs have the following properties. Each class consists of *core nodes* and *satellite nodes*. Core nodes of the same class are tightly connected with each other and do not have any erroneous edges. Satellite nodes are relatively weakly connected to core nodes of the same class. The satellite nodes are also connected to some other classes’ satellite nodes (i.e., introducing errors). This core-satellite model is intended to simulate real-world data in which some data points are close to the class boundaries (satellite nodes). More precisely, graphs 6–10 were generated as follows. Each graph consists of 2000 nodes ($m = 2000$) uniformly distributed over 10 classes ($K = 10$). 10% of the nodes are the core nodes. For every core node, we randomly choose 10 other core nodes of the same class and connect it to them with edge weight 1 (that is, each core node is connected to at least 10 core nodes of the same class). For every satellite node, we randomly choose one core node of the same class and connect them with edge weight 0.01. Also, for each satellite node, we randomly choose one satellite node of some other class (i.e., introducing error) and connect them with edge weight w_e . We set the error edge weight $w_e = 0.002, 0.004, \dots, 0.01$ for graphs 6, 7, \dots , 10, respectively. Note that although classes are uniformly distributed, pure components that optimize the generalization bound may be non-uniform in size. For graphs generated in this manner, degrees vary within the same class since the satellite nodes have smaller degrees than the core nodes. Our analysis suggests that \mathbf{L} -scaling will do poorly. Figure 3 shows that on the five core-satellite graphs, \mathbf{K} -scaling indeed produces higher performance than \mathbf{L} -scaling. In particular, \mathbf{K} -scaling does well even when \mathbf{L} -scaling rather underperforms the unnormalized baseline.

Our analysis suggests that \mathbf{K} -scaling should work well when the graph has relatively small error. This trend is more clearly observed on these core-satellite graphs without dimension reduction. As shown in Figure 4, the advantage of \mathbf{K} -scaling over \mathbf{L} -scaling is more prominent on the graphs with smaller error edge weights. On the other hand, the theory suggests that when the graph has large error (large cut), the benefit of normalization is less clear (since the derivation of \mathbf{K} -scaling assumes near-zero cut). This is especially so when dimension reduction is not applied because as pointed out in Section 3.5, dimension reduction reduces error. This trend can be observed in Figure 5, which shows that on graphs 1–3 (having larger errors than the core-satellite graphs), neither \mathbf{L} -scaling nor \mathbf{K} -scaling prominently improves performance over the unnormalized Laplacian without dimension reduction though \mathbf{L} -scaling seems to perform slightly better. Note that the performance without dimension reduction (Figure 5) is significantly worse than the performance with dimension reduction (Figure 2). This means that dimension reduction, which reduces error, is important when we try to apply graph based methods.

Additionally, we show illustrative toy examples based on the core-satellite model. Given the original graph as in Figure 6 (a), Figure 6 (b)–(d) show the graphs corresponding to the scaled adjacency matrices $\mathbf{S}^{-1/2}\mathbf{W}\mathbf{S}^{-1/2}$ where \mathbf{S} is derived from \mathbf{L} -scaling, \mathbf{K} -scaling with $\alpha = 0.01$, and \mathbf{K} -scaling with $\alpha = 0.1$, respectively. We observe that, compared with the unnormalized case, \mathbf{K} -scaling and \mathbf{L} -scaling essentially balance the edge weights between the two classes (i.e., “normalizing”) by relatively lowering the weights of class2 which is more “massive”. However, in this example, \mathbf{L} -scaling in a sense overdoes it and so is rather harmful as it amplifies the error edge weights over the weights of the within-class edges. \mathbf{K} -scaling does not suffer from this problem.




Figure 4: Classification accuracy on the core-satellite graphs. x -axis: error edge weight w_e . $n = 40, m = 2000$. Without dimension reduction. Average over 10 random splits.




Figure 5: Classification accuracy (%) on the graphs where degrees are nearly constant within the class. Average over 10 random splits. $n = 40, m = 2000$. Without dimension reduction.

4.3 Real-world Data Experiments

Our real-world data experiments use two image data sets (MNIST and UMIST) and one text data set (RCV1).

4.3.1 DATA AND BASELINE

The MNIST data set, downloadable from <http://yann.lecun.com/exdb/mnist/>, consists of hand-written digit image data (representing 10 classes, from digit “0” to “9”). For our experiments, we randomly choose 2000 images (i.e., $m = 2000$). The UMIST data set, downloadable from <http://images.ee.umist.ac.uk/danny/database.html>, consists of 575 face images taken from several angles of 20 people (representing 20 classes). The details of this data are described in Graham and Allinson (1998). We use all the images (i.e., $m = 575$). Reuters Corpus Version 1 (RCV1) consists of news articles labeled with topics. For our experiments, we chose 10 topics (representing 10 classes) that have relatively large populations and randomly chose 2000 articles that are labeled with exactly one of those 10 topics. The class distribution over these 2000 articles is non-uniform as shown in Figure 7.

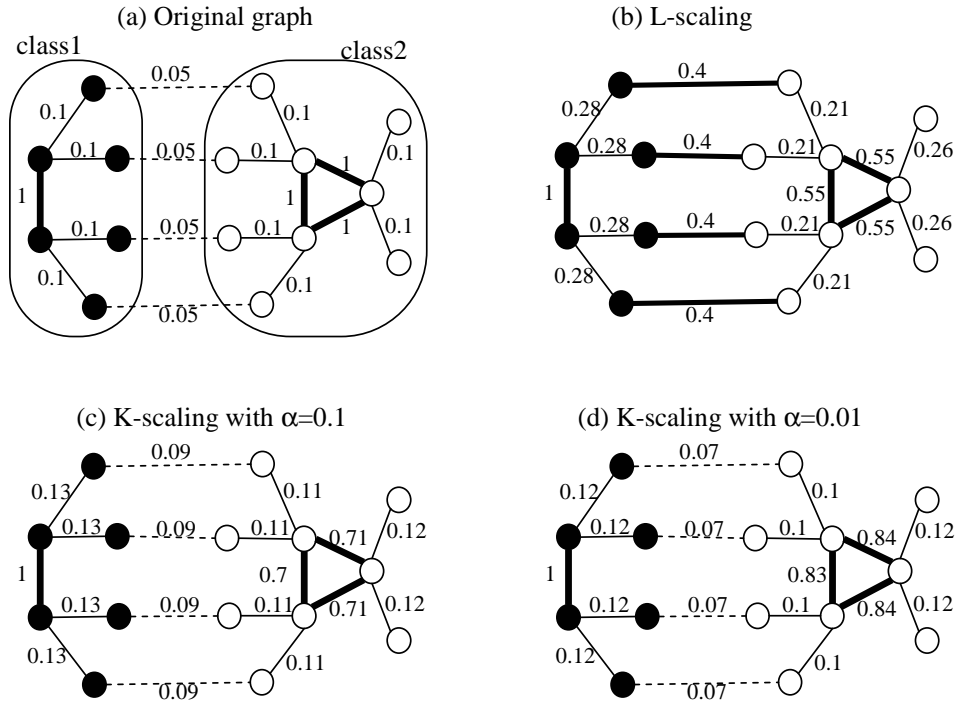


Figure 6: Illustrative toy examples of scaled adjacency matrices $\mathbf{S}^{-1/2}\mathbf{W}\mathbf{S}^{-1/2}$ where \mathbf{S} is derived from \mathbf{L} -scaling or \mathbf{K} -scaling. For the two-class core-satellite graph in (a), \mathbf{L} -scaling makes the weights of error edges larger than the edge weights between the core nodes and satellite nodes of the same class as in (b). \mathbf{K} -scaling does not suffer from this problem ((c),(d)). (For an easy comparison, in (b)–(d), edge weights are multiplied with constants so that the largest weight becomes one.)


To generate graphs from the image data, as is commonly done, we first generate the vectors of the gray-scale values of the pixels, and produce the edge weight between the i -th and the j -th data points \mathbf{X}_i and \mathbf{X}_j by $w_{i,j} = \exp(-\|\mathbf{X}_i - \mathbf{X}_j\|^2/t)$ where $t > 0$ is a parameter (radial basis function (RBF) kernels). To generate graphs from the text data, we first create the bag-of-words vectors using content words only² and then set $w_{i,j}$ based on RBF as above or set $w_{i,j}$ to the inner product of \mathbf{X}_i and \mathbf{X}_j (linear kernels). Optionally, we zero out all $w_{i,j}$ but k nearest neighbors (i.e., i is j 's k nearest neighbors or j is i 's k nearest neighbors) to reduce error in graphs and refer to it as the RBF (or linear) kernel with k NN.

As our baseline, we also test the supervised configuration by letting $\mathbf{W} + \beta\mathbf{I}$ (where \mathbf{W} is a weight matrix whose (i, j) -entry is $w_{i,j}$) be the kernel matrix and using the same least squares loss function. We set β to the optimum, which was 0.001 for the RBF kernel for RCV1 and 1 for the other graphs.

2. To generate a bag-of-words vector from a document, as is commonly done, we remove function words (such as “a”, “the”, and so on), set word frequencies of the document to the corresponding vector entries, and then scale the vector into a unit vector.

GPOL	Domestic politics	486
GSPO	Sports	407
GDIP	International relations	299
GCRIM	Crime, law enforcement	224
GJOB	Labor issues	206
GVIO	War, civil war	142
GDIS	Disasters and accidents	89
GHEA	Health	57
GENT	Arts, culture, entertainment	47
GENV	Environments	43
Total		2000

Figure 7: 10 RCV1 categories and their populations used in our experiments.


 Figure 8: Classification accuracy (%) in relation to the number of labeled examples (n) on MNIST. $m = 2000$. (a) Dimensionality and α were determined by cross validation. (b) Dimensionality and α were set to the optimum. Average over 10 random splits.

4.3.2 RESULTS

Figure 8 shows performance in relation to the number of labeled examples (n) on the MNIST data set. The comparison of the three bold lines (representing the methods with dimension reduction) in Figure 8 (a) shows that when the dimensionality and α are determined (performing simple 2-dimensional grid search) by cross validation, **K**-scaling outperforms **L**-scaling, and **L**-scaling outperforms the unnormalized Laplacian. These performance differences are statistically significant ($p \leq 0.01$) based on the paired t test. The performance of the unnormalized Laplacian (with dimension reduction) is roughly consistent with the performance with similar (m, n) with heuristic dimension selection in Belkin and Niyogi (2004). Although without dimension reduction, **L**-scaling

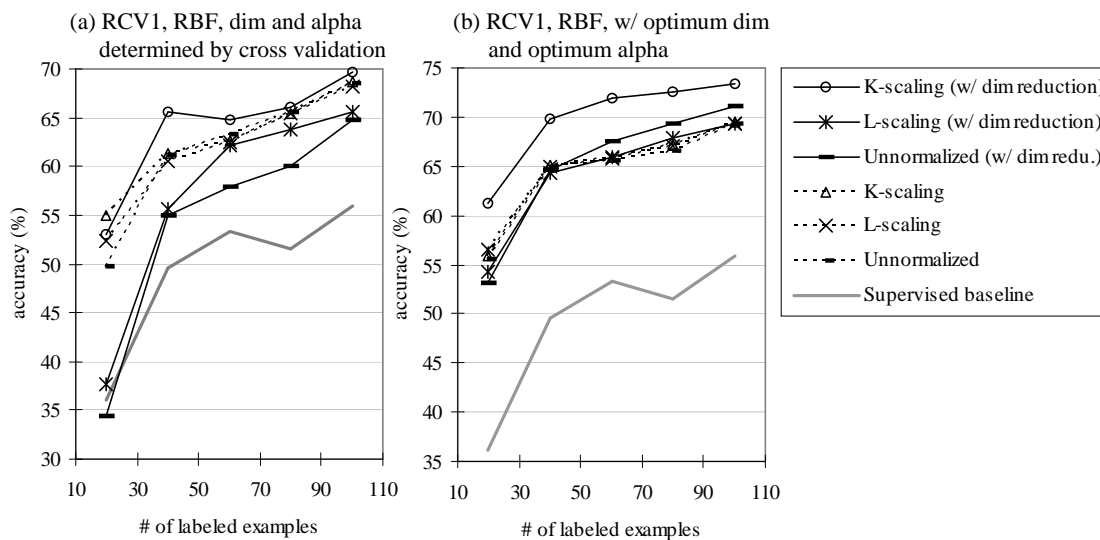


Figure 9: Classification accuracy (%) in relation to the number of labeled examples (n) on RCV1. RBF kernel (with $t = 0.25$). $m = 2000$. (a) Dimensionality and α were determined by cross validation. (b) Dimensionality and α were set to the optimum. Performance differences of the best performing method ‘K-scaling (w/ dim reduction)’ from ‘L-scaling (w/ dim reduction)’ and ‘Unnormalized (w/ dim redu.)’ are statistically significant ($p \leq 0.01$) in both the settings (a) and (b).

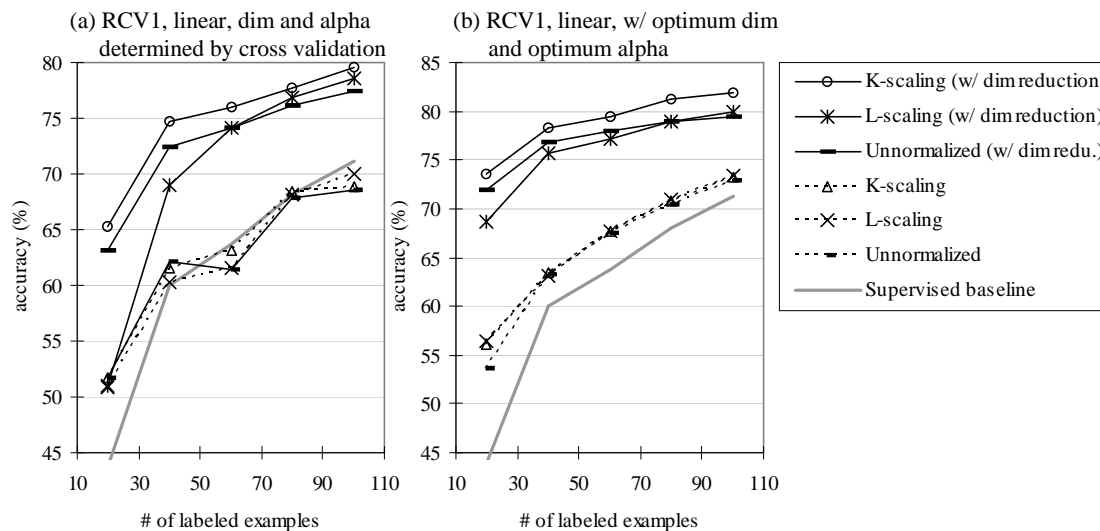


Figure 10: Classification accuracy (%) in relation to the number of labeled examples (n) on RCV1. Linear kernel. $m = 2000$. (a) Dimensionality and α were determined by cross validation. (b) Dimensionality and α were set to the optimum. Performance differences of the best performing method ‘K-scaling (w/ dim reduction)’ from the second and third best ‘L-scaling (w/ dim reduction)’ and ‘Unnormalized (w/ dim redu.)’ are statistically significant ($p \leq 0.01$) in both the settings (a) and (b).

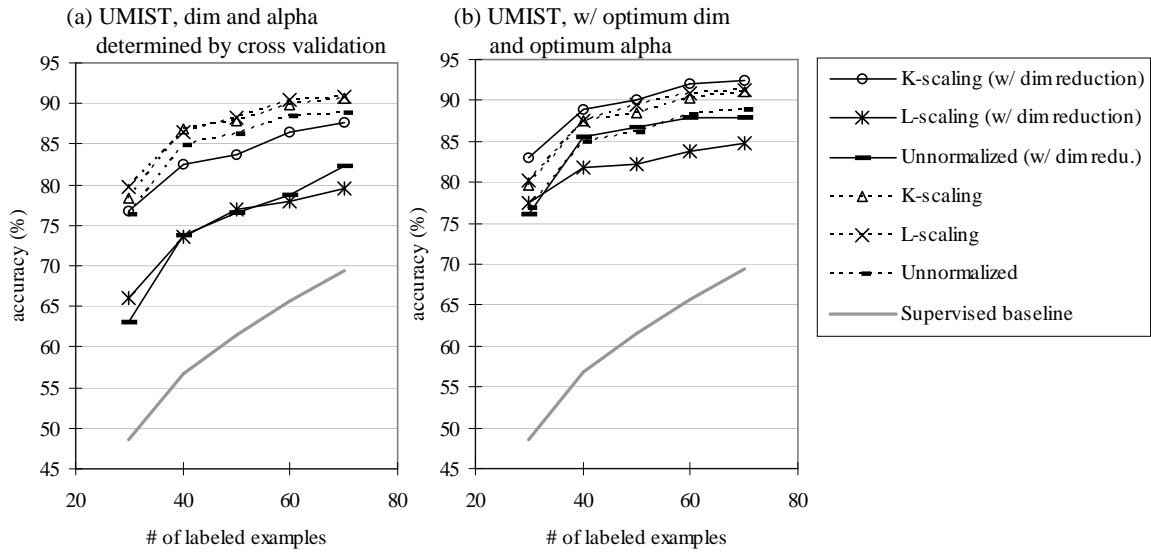


Figure 11: Classification accuracy (%) in relation to the number of labeled examples (n) on UMIST. $m = 575$. (a) Dimensionality and α were determined by cross validation. (b) Dimensionality and α were set to the optimum. In (b), performance differences of the best performing method ‘**K**-scaling (w/ dim reduction)’ from the second and third best ‘**K**-scaling’ and ‘**L**-scaling’ are statistically significant ($p \leq 0.01$).

and **K**-scaling still improve performance over the unnormalized Laplacian, the best performance is always obtained by **K**-scaling with dimension reduction (the bold line with circles).

In Figure 8 (a), the unnormalized Laplacian with dimension reduction underperforms the unnormalized Laplacian without dimension reduction, indicating that dimension reduction rather degrades performance in this case. By comparing Figure 8 (a) and (b), we observe that this seemingly counter-intuitive performance trend is caused by the difficulty in choosing the right dimensionality by cross validation. Figure 8 (b) shows the performance at the optimum dimensionality and the optimum α . As observed, if the optimum dimensionality is known (as in (b)), dimension reduction improves performance either with or without normalization by **K**-scaling and **L**-scaling, and that all the transductive configurations outperform the supervised baseline. We also note that the comparison of Figure 8 (a) and (b) shows that choosing good dimensionality by cross validation is much harder than choosing α by cross validation especially when the number of labeled examples is small. This trend was observed also on the other data sets on which we experimented.

On the RCV1 data set, the performance trend is essentially similar to that of MNIST. Figure 9 shows the performance on RCV1 using the RBF kernel ($t = 0.25$, 100NN). In the setting of Figure 9 (a) where the dimensionality and α were determined by cross validation, **K**-scaling with dimension reduction generally performs the best. By setting the dimensionality and α to the optimum, the benefit of **K**-scaling with dimension reduction is even clearer (Figure 9 (b)).

On text data like RCV1, linear kernels (instead of RBF) are often used. Figure 10 shows the performance with linear kernels with 100NN. Again, **K**-scaling with dimension reduction performs the best. Its performance differences from the second and third best ‘**L**-scaling (w/ dim reduction)’

and ‘Unnormalized (w/ dim redu.)’ are statistically significant ($p \leq 0.01$) in both Figure 10 (a) and (b).

In Figure 11, we observe that dimension reduction seems less useful on the UMIST data set. We conjecture that this may be because UMIST differs from our other data sets in that it is much more ‘sparse’; UMIST has a smaller number of data points ($m = 575$ vs. $m = 2000$) while it has more classes ($K = 20$ vs. $K = 10$). Nevertheless, when the dimensionality and α are set to the optimum (Figure 11 (b)), again, **K**-scaling with dimension reduction performs the best. Its differences from the second and the third best methods (**K**-scaling without dimension reduction and **L**-scaling without dimension reduction) are statistically significant ($p \leq 0.01$).

Overall, on these graphs generated from image and text data sets, **K**-scaling with dimension reduction consistently outperformed the others. But without dimension reduction, **K**-scaling and **L**-scaling were not always effective. Transductive learning (either with or without normalization) generally improved performance.

4.4 Approximation of **K**-scaling

Although **K**-scaling consistently improves performance as shown above, its drawback is the relatively large runtime as it involves the computation of the inverse of an m -by- m matrix. We propose a less computationally-intensive approximation method using a known fact that $(\mathbf{I} - \mathbf{A})^{-1} = \sum_{k=0}^{\infty} \mathbf{A}^k$ if $\|\mathbf{A}\|_2 < 1$. As in the introduction, let $\mathbf{D} = \text{diag}_i(\text{deg}_i(G))$, and let \mathbf{W} be a weight matrix such that $\mathbf{W}_{i,j} = w_{i,j}$ so that we can write $\mathcal{L} = \mathbf{D} - \mathbf{W}$. Let $\hat{\mathbf{D}} = \mathbf{D} + \alpha\mathbf{I}$. We define $\hat{\mathbf{K}}(h)$ to be the h -th order approximation of $\bar{\mathbf{K}} = (\mathcal{L} + \alpha\mathbf{I})^{-1}$ as follows:

$$\hat{\mathbf{K}}(h) = \hat{\mathbf{D}}^{-1/2} \left(\sum_{k=0}^h \left(\hat{\mathbf{D}}^{-1/2} \mathbf{W} \hat{\mathbf{D}}^{-1/2} \right)^k \right) \hat{\mathbf{D}}^{-1/2} .$$

We then set the i -th scaling factor \mathbf{S}_i so that:

$$\mathbf{S}_i = \hat{\mathbf{K}}(h)_{i,i}^{-1} .$$

Since $\lim_{h \rightarrow \infty} \hat{\mathbf{K}}(h) = \bar{\mathbf{K}}$, the scaling factors produced with a sufficiently large h closely approximate **K**-scaling. On the other hand, since $\hat{\mathbf{K}}(0) = \hat{\mathbf{D}}^{-1} = (\mathbf{D} + \alpha\mathbf{I})^{-1}$, the scaling factors produced by $\hat{\mathbf{K}}(0)$ with $\alpha = 0$ are exactly the same as **L**-scaling (or the standard degree-scaling).

Figure 12 shows the performance of this approximation method with $h = 0, 2, 5, 10$ with dimension reduction in comparison with corresponding **K**-scaling and **L**-scaling on MNIST. In Figure 12 (b), we observe that at the optimum dimensionality and α , the performance of the approximation method lies exactly between that of **L**-scaling and **K**-scaling, and it approaches to **K**-scaling as the order h increases. Intuitively, with a larger h , this approximation method takes more and more global connections into account and improves performance.

5. Conclusion

We derived generalization bounds for multi-category classification on graphs with Laplacian regularization, using geometric properties of the graph. In particular, we used this analysis to obtain a better understanding of the role of normalization of the graph Laplacian matrix. We argued that the standard **L**-scaling normalization method has the undesirable property that the normalization factors

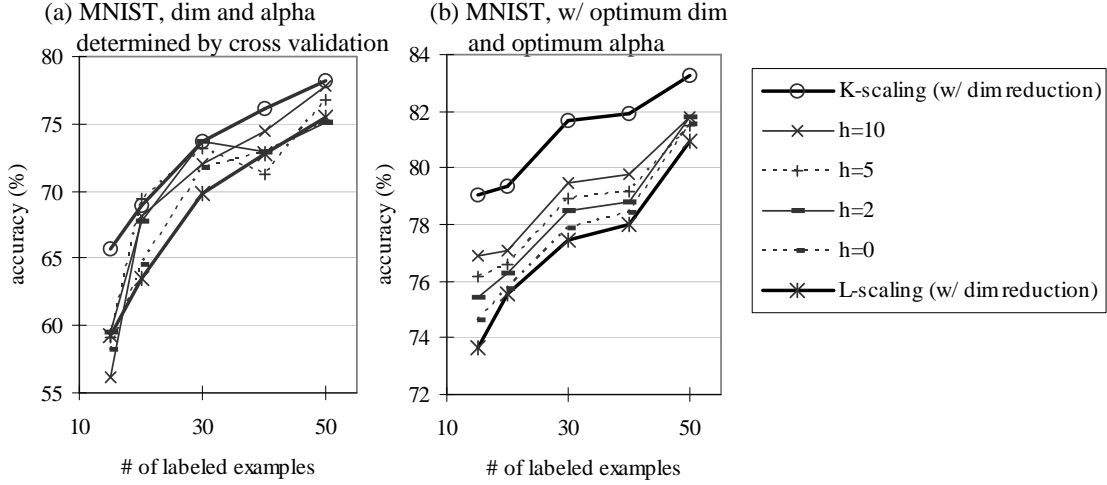


Figure 12: Classification accuracy (%) of the approximation method using $\hat{\mathbf{K}}(h)$. MNIST. (a) Dimensionality and α were determined by cross validation. (b) Dimensionality and α were set to the optimum.

can vary significantly within a pure component. An alternate normalization method, which we call **K-scaling**, is proposed to remedy the problem. Experiments confirm the superiority of **K-scaling** combined with dimension reduction. Finally, there are possible extensions of this work that require further investigation, for example, how to use the **K-scaling** for other types of graphs such as directed graphs, and how to apply this idea to spectral clustering.

Appendix A. Proof of Theorem 1

The proof employs the stability analysis of Zhang (2003), and is similar to the proof of a related bound for binary-classification in Zhang and Ando (2006). We shall introduce the following notation. let $i_{n+1} \neq i_1, \dots, i_n$ be an integer randomly drawn from \bar{Z}_n , and let $Z_{n+1} = Z_n \cup \{i_{n+1}\}$. Let $\hat{\mathbf{f}}(Z_{n+1})$ be the semi-supervised learning method (1) using training data in Z_{n+1} :

$$\hat{\mathbf{f}}(Z_{n+1}) = \arg \min_{\mathbf{f} \in \mathbb{R}^{mK}} \left[\frac{1}{n} \sum_{j \in Z_{n+1}} \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q}_K \mathbf{f} \right].$$

We have the following stability lemma (a related result can be found in Zhang, 2003);

Lemma 8 *The following inequality holds for each $k = 1, \dots, K$:*

$$|\hat{f}_{i_{n+1},k}(Z_{n+1}) - \hat{f}_{i_{n+1},k}(Z_n)| \leq |\nabla_{1,k} \phi(\hat{\mathbf{f}}_{i_{n+1}}(Z_{n+1}), y_{i_{n+1}})| \mathbf{K}_{i_{n+1},i_{n+1}} / (2\lambda n),$$

where $\nabla_{1,k} \phi(f_i, y)$ denotes a sub-gradient of $\phi(f_i, y)$ with respect to $f_{i,k}$, where $\mathbf{f}_i = [f_{i,1}, \dots, f_{i,K}]$.

Proof From Rockafellar (1970), we know that there exist sub-gradients $\nabla_{1,k} \phi$ such that the following first-order condition for the optimization problem (1) holds:

$$-2\lambda n \mathbf{K}^{-1} \hat{\mathbf{f}}_{\cdot,k}(Z_n) = \sum_{j \in Z_n} \nabla_{1,k} \phi(\hat{\mathbf{f}}_j(Z_n), y_j) \mathbf{e}_j,$$

where \mathbf{e}_j is the m -dimensional vector with all zeros except for the j -component with value one. Similarly, we have

$$-2\lambda n \mathbf{K}^{-1} \hat{\mathbf{f}}_{\cdot,k}(Z_{n+1}) = \sum_{j \in Z_{n+1}} \nabla_{1,k} \phi(\hat{\mathbf{f}}_j(Z_{n+1}), y_j) \mathbf{e}_j.$$

Now, for simplicity, let $\mathbf{g} = \hat{\mathbf{f}}(Z_n)$ and $\mathbf{h} = \hat{\mathbf{f}}(Z_{n+1})$. By subtracting the above two equations, and then taking the inner product with $\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k}$, we obtain

$$\begin{aligned} -2\lambda n (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k})^T \mathbf{K}^{-1} (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k}) &= \nabla_{1,k} \phi(h_{i_{n+1}}, y_{i_{n+1}}) (h_{i_{n+1},k} - g_{i_{n+1},k}) \\ &\quad + \sum_{j \in Z_n} (\nabla_{1,k} \phi(\mathbf{h}_j, y_j) - \nabla_{1,k} \phi(\mathbf{g}_j, y_j)) (h_{j,k} - g_{j,k}). \end{aligned}$$

Note that if $c(s)$ is a convex function of s , then it is easy to verify that $(\nabla c(s_1) - \nabla c(s_2))(s_1 - s_2) \geq 0$. Therefore we have $\sum_{j \in Z_n} (\nabla_{1,k} \phi(\mathbf{h}_j, y_j) - \nabla_{1,k} \phi(\mathbf{g}_j, y_j)) (h_{j,k} - g_{j,k}) \geq 0$. This implies that

$$2\lambda n (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k})^T \mathbf{K}^{-1} (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k}) \leq -\nabla_{1,k} \phi(\mathbf{h}_{i_{n+1}}, y_{i_{n+1}}) (h_{i_{n+1},k} - g_{i_{n+1},k}).$$

Using Cauchy-Schwartz inequality, we have

$$\begin{aligned} 2\lambda n (h_{i_{n+1},k} - g_{i_{n+1},k})^2 &= 2\lambda n ((\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k})^T \mathbf{e}_{i_{n+1}})^2 \\ &\leq 2\lambda n (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k})^T \mathbf{K}^{-1} (\mathbf{h}_{\cdot,k} - \mathbf{g}_{\cdot,k}) \mathbf{e}_{i_{n+1}}^T \mathbf{K} \mathbf{e}_{i_{n+1}} \\ &\leq |\nabla_{1,k} \phi(h_{i_{n+1}}, y_{i_{n+1}})| \cdot |h_{i_{n+1},k} - g_{i_{n+1},k}| \mathbf{K}_{i_{n+1},i_{n+1}}. \end{aligned}$$

Therefore we have $|h_{i_{n+1},k} - g_{i_{n+1},k}| \leq |\nabla_{1,k} \phi(h_{i_{n+1}}, y_{i_{n+1}})| \mathbf{K}_{i_{n+1},i_{n+1}} / (2\lambda n)$. ■

Lemma 9 *The following inequality holds*

$$\mathbf{err}(\hat{\mathbf{f}}_{i_{n+1}}(Z_n), y_{i_{n+1}}) \leq \sup_{k=k_0, i_{n+1}} \left[\frac{1}{a} \phi_0(\hat{f}_{i_{n+1},k}(Z_{n+1}), \delta_{i_{n+1},k}) + \left(\frac{b}{c\lambda n} \mathbf{K}_{i_{n+1},i_{n+1}} \right)^p \right].$$

Proof If $\hat{\mathbf{f}}(Z_n)$ does not make an error on the i_{n+1} -th example. That is, if $\mathbf{err}(\hat{\mathbf{f}}_{i_{n+1}}(Z_n), y_{i_{n+1}}) = 0$, then the inequality automatically holds.

Now, assume that $\hat{\mathbf{f}}(Z_n)$ makes an error on the i_{n+1} -th example: $\mathbf{err}(\hat{\mathbf{f}}_{i_{n+1}}(Z_n), y_{i_{n+1}}) = 1$. Then there exists $k_0 \neq y_{i_{n+1}}$ such that $\hat{f}_{i_{n+1},y_{i_{n+1}}}(Z_n) \leq \hat{f}_{i_{n+1},k_0}(Z_n)$. This means that for any d , either $\hat{f}_{i_{n+1},y_{i_{n+1}}}(Z_n) \leq d$ or $\hat{f}_{i_{n+1},k_0}(Z_n) \geq d$. We simple let $d = (\inf\{x : \phi_0(x, 1) \leq a\} + \sup\{x : \phi_0(x, 0) \leq a\})/2$. By the definition of c , either we have $\inf\{x : \phi_0(x, 1) \leq a\} - \hat{f}_{i_{n+1},y_{i_{n+1}}}(Z_n) \geq c/2$ or we have $\hat{f}_{i_{n+1},k_0}(Z_n) - \sup\{x : \phi_0(x, 0) \leq a\} \geq c/2$. It follows that there exists $k = k_0$ or $k = y_{i_{n+1}}$ such that either $\phi_0(\hat{f}_{i_{n+1},k}(Z_{n+1}), \delta_{y_{i_{n+1}},k}) \geq a$ or $|\hat{f}_{i_{n+1},k}(Z_{n+1}) - \hat{f}_{i_{n+1},k}(Z_n)| \geq c/2$.

Using Lemma 8, we have either $\phi_0(\hat{f}_{i_{n+1},k}(Z_{n+1}), \delta_{y_{i_{n+1}},k}) \geq a$ or $b\mathbf{K}_{i_{n+1},i_{n+1}} / (2\lambda n) \geq c/2$, implying that

$$\frac{1}{a} \phi_0(\hat{f}_{i_{n+1},k}(Z_{n+1}), \delta_{y_{i_{n+1}},k}) + \left(\frac{b\mathbf{K}_{i_{n+1},i_{n+1}}}{c\lambda n} \right)^p \geq 1 = \mathbf{err}(\hat{\mathbf{f}}_{i_{n+1}}(Z_n), y_{i_{n+1}}).$$
■

We are now ready to prove Theorem 1. For every $j \in Z_{n+1}$, denote by $Z_{n+1}^{(j)}$ the subset of n samples in Z_{n+1} with the j -th data point left out. From Lemma 9, we have

$$\mathbf{err}(\hat{\mathbf{f}}_j(Z_{n+1}^{(j)}), y_j) \leq \frac{1}{a} \phi(\hat{\mathbf{f}}_j(Z_{n+1}), y_j) + \left(\frac{b}{c\lambda n} \mathbf{K}_{j,j} \right)^p.$$

We thus obtain for all $\mathbf{f} \in R^{mK}$:

$$\begin{aligned} \sum_{j \in Z_{n+1}} \mathbf{err}(\hat{\mathbf{f}}_j(Z_{n+1}^{(j)}), y_j) &\leq \frac{1}{a} \sum_{j \in Z_{n+1}} \phi(\hat{\mathbf{f}}_j(Z_{n+1}), y_j) + \sum_{j \in Z_{n+1}} \left(\frac{b}{c\lambda n} \mathbf{K}_{j,j} \right)^p \\ &\leq \frac{1}{a} \left[\sum_{j \in Z_{n+1}} \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q} \mathbf{K} \mathbf{f} \right] + \sum_{j \in Z_{n+1}} \left(\frac{b}{c\lambda n} \mathbf{K}_{j,j} \right)^p. \end{aligned}$$

Therefore

$$\begin{aligned} &\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in Z_n} \mathbf{err}(\hat{\mathbf{f}}_j(Z_n), y_j) \\ &= \mathbf{E}_{Z_{n+1}} \mathbf{err}(\hat{\mathbf{f}}_{n+1}(Z_{n+1}^{(n+1)}), y_{n+1}) = \frac{1}{n+1} \mathbf{E}_{Z_{n+1}} \sum_{j \in Z_{n+1}} \mathbf{err}(\hat{\mathbf{f}}_j(Z_{n+1}^{(j)}), y_j) \\ &\leq \frac{n}{a(n+1)} \mathbf{E}_{Z_{n+1}} \left[\frac{1}{n} \sum_{j \in Z_{n+1}} \phi(\mathbf{f}_j, y_j) + \lambda \mathbf{f}^T \mathbf{Q} \mathbf{K} \mathbf{f} \right] + \frac{1}{n+1} \mathbf{E}_{Z_{n+1}} \sum_{j \in Z_{n+1}} \left(\frac{b}{c\lambda n} \mathbf{K}_{j,j} \right)^p \\ &= \frac{1}{a} \left[\frac{1}{m} \sum_{j=1}^m \phi(\mathbf{f}_j, y_j) + \frac{\lambda n}{n+1} \mathbf{f}^T \mathbf{Q} \mathbf{K} \mathbf{f} \right] + \frac{1}{m} \sum_{j=1}^m \left(\frac{b \mathbf{K}_{j,j}}{c\lambda n} \right)^p. \end{aligned}$$

Appendix B. Proof of Theorem 7

The idea is similar to that of Theorem 5. We use the same notation, and let \mathbf{v}_ℓ be the indicator function of V_ℓ in V . Let \mathbf{I}_ℓ be the diagonal matrix with value ones for nodes corresponding to V_ℓ and zeros elsewhere. We have $\forall \mathbf{u} = [u_1, \dots, u_m] \in R^m$, $(\mathbf{u} - \mathbf{u}^T \mathbf{v}_\ell \mathbf{v}_\ell / m_\ell)^T \mathbf{v}_\ell = 0$. Therefore by the definition of $\lambda_2(G_\ell)$, we have

$$\begin{aligned} \frac{1}{2} \sum_{j, j' \in V_\ell} w_{j,j'} (u_j - u_{j'})^2 &= (\mathbf{I}_\ell \mathbf{u} - \mathbf{u}^T \mathbf{v}_\ell \mathbf{v}_\ell / m_\ell)^T \mathcal{L}(G_\ell) (\mathbf{I}_\ell \mathbf{u} - \mathbf{u}^T \mathbf{v}_\ell \mathbf{v}_\ell / m_\ell) \\ &\geq \lambda_2(G_\ell) \|\mathbf{I}_\ell \mathbf{u} - \mathbf{u}^T \mathbf{v}_\ell \mathbf{v}_\ell / m_\ell\|_2^2 = \lambda_2(G_\ell) \mathbf{u}^T [\mathbf{I}_\ell - \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell] \mathbf{u}. \end{aligned}$$

We thus obtain

$$\begin{aligned} \mathbf{u}^T \mathcal{L} \mathbf{u} &= \frac{1}{2} \sum_{j, j'=1}^m w_{j,j'} (u_j - u_{j'})^2 \geq \sum_{\ell=1}^q \frac{1}{2} \sum_{j, j' \in V_\ell} w_{j,j'} (u_j - u_{j'})^2 \\ &\geq \mathbf{u}^T \left[\sum_{\ell=1}^q \lambda_2(G_\ell) (\mathbf{I}_\ell - \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell) \right] \mathbf{u}. \end{aligned}$$

Therefore $\mathcal{L} - \sum_{\ell=1}^q \lambda_2(G_\ell)(\mathbf{I}_\ell - \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell)$ is positive semi-definite, and thus

$$\left(\alpha \mathbf{I} + \sum_{\ell=1}^q \lambda_2(G_\ell)(\mathbf{I}_\ell - \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell) \right)^{-1} - (\alpha \mathbf{I} + \mathcal{L})^{-1}$$

is positive-semi-definite. Also, from (2) we have $\mathbf{S}^{-1/2} \mathbf{K} \mathbf{S}^{-1/2} = (\alpha \mathbf{I} + \mathcal{L}(G))^{-1}$, so we know that the diagonal entries of $\mathbf{S}^{-1/2} \mathbf{K} \mathbf{S}^{-1/2}$ can be upper-bounded by those of

$$\left(\alpha \mathbf{I} + \sum_{\ell=1}^q \lambda_2(G_\ell)(\mathbf{I}_\ell - \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell) \right)^{-1} = \sum_{\ell=1}^q (\alpha + \lambda_2(G_\ell))^{-1} (\mathbf{I}_\ell + \alpha^{-1} \lambda_2(G_\ell) \mathbf{v}_\ell \mathbf{v}_\ell^T / m_\ell).$$

For the latter, its m_ℓ diagonal entries for each pure component ℓ can be upper bounded by $\lambda_2(G_\ell)^{-1} + (\alpha m_\ell)^{-1}$. Therefore:

$$\begin{aligned} m^{1/p} \mathbf{tr}_p(\mathbf{K}) &\leq \left(\sum_{\ell=1}^q s_\ell(p) (\alpha^{-1} m_\ell^{-1} + \lambda_2(G_\ell)^{-1})^p \right)^{1/p} \\ &\leq \left[\alpha^{-1} \left(\sum_{\ell=1}^q s_\ell(p) m_\ell^{-p} \right)^{1/p} + \left(\sum_{\ell=1}^q s_\ell(p) \lambda_2(G_\ell)^{-p} \right)^{1/p} \right]. \end{aligned}$$

Substitute this estimate into Theorem 4, and we have

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \tilde{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j(Z_n), y_j) \leq \frac{C_p(a, b, c)}{n^{p/(p+1)}} \left[m^{-1/p} (\alpha s + C) (\alpha^{-1} A + B) \right]^{p/(p+1)},$$

where $A = (\sum_{\ell=1}^q s_\ell(p) m_\ell^{-p})^{1/p}$, $B = (\sum_{\ell=1}^q s_\ell(p) \lambda_2(G_\ell)^{-p})^{1/p}$, and $C = \mathbf{cut}(\mathcal{L}_S, y)$. Now optimize over α (let $\alpha = \sqrt{AC/(sB)}$), and simplify to obtain the desired inequality.

Appendix C. Dimension Reduction

As in Section 3.5, denote by $\mathbf{P}_S^r(G)$ the projection operator onto the eigenspace of $\alpha \mathbf{S}^{-1} + \mathcal{L}_S(G)$ corresponding to the r smallest eigenvalues, and define the following regularizer on the reduced subspace:

$$\mathbf{f}_{\cdot, k}^T \mathbf{K}^{-1} \mathbf{f}_{\cdot, k} = \begin{cases} \mathbf{f}_{\cdot, k}^T \mathbf{K}_0^{-1} \mathbf{f}_{\cdot, k} & \text{if } \mathbf{P}_S^r(G) \mathbf{f}_{\cdot, k} = \mathbf{f}_{\cdot, k}, \\ +\infty & \text{otherwise.} \end{cases}$$

Note that in the following, we will focus on bounding the generalization complexity using the reduced dimensionality r . In such context, the choice of \mathbf{K}_0 is not important as far as our analysis is concerned. We may simply choose $\mathbf{K}_0 = \mathbf{I}$ (or we may let $\mathbf{K}_0^{-1} = \alpha \mathbf{S}^{-1} + \mathcal{L}_S(G)$).

The following theorem shows that the target vectors can be well approximated by their projection via $\mathbf{P}_S^q(G)$.

Theorem 10 *Let $G' = \cup_{\ell=1}^q G_\ell$ ($G_\ell = (V_\ell, E_\ell)$) be a pure subgraph of G . Consider $r \geq q$. Then we have:*

$$\lambda_{r+1}(\mathcal{L}_S(G)) \geq \lambda_{r+1}(\mathcal{L}_S(G')) \geq \min_{\ell} \lambda_2(\mathcal{L}_S(G_\ell)).$$

For each k , let $\bar{f}_{j,k} = \delta_{y_j,k}$ be the target (encoding of the true labels) for class k ($j = 1, \dots, m$). Then $\|\mathbf{P}_S^r(G)\bar{\mathbf{f}}_{\cdot,k} - \bar{\mathbf{f}}_{\cdot,k}\|_2^2 \leq \delta_r(\mathbf{S})\|\bar{\mathbf{f}}_{\cdot,k}\|_2^2$, where

$$\delta_r(\mathbf{S}) = \frac{\|\mathcal{L}_S(G) - \mathcal{L}_S(G')\|_2 + d(\mathbf{S})}{\lambda_{r+1}(\mathcal{L}_S(G))}, \quad d(\mathbf{S}) = \max_{\ell} \frac{1}{2|V_{\ell}|} \sum_{j,j' \in V_{\ell}} w_{j,j'} (\mathbf{S}_j^{-1/2} - \mathbf{S}_{j'}^{-1/2})^2.$$

Proof Let $\mathbf{E} = \mathcal{L}_S(G) - \mathcal{L}_S(G')$, then \mathbf{E} is a positive semi-definite matrix. Therefore, we know that the eigenvalues of $\mathcal{L}_S(G)$ are no less than the corresponding eigenvalues of $\mathcal{L}_S(G')$. This follows easily from the minimax variational formulation for the $q+1$ -th smallest eigenvalue of a matrix $\mathcal{L} \in \mathbb{R}^{m \times m}$, which can be written as

$$\max_{V_q \in \mathbb{R}^{m \times q}: V_q^T V_q = I_{q \times q}} \min_{\mathbf{v} \in \mathbb{R}^m: \|\mathbf{v}\|_2=1, V_q^T \mathbf{v}=0} \mathbf{v}^T \mathcal{L} \mathbf{v},$$

together with the fact that $\mathbf{v}^T \mathcal{L}_S(G) \mathbf{v} \geq \mathbf{v}^T \mathcal{L}_S(G') \mathbf{v}$.

Since the $(q+1)$ -th smallest eigenvalue of $\mathcal{L}_S(G')$ is $\min_{\ell} \lambda_2(\mathcal{L}(G_{\ell}))$, we obtain the first displayed inequalities. Moreover,

$$\bar{\mathbf{f}}_{\cdot,k}^T \mathcal{L}_S(G) \bar{\mathbf{f}}_{\cdot,k} = \bar{\mathbf{f}}_{\cdot,k}^T \mathbf{E} \bar{\mathbf{f}}_{\cdot,k} + \bar{\mathbf{f}}_{\cdot,k}^T \mathcal{L}_S(G') \bar{\mathbf{f}}_{\cdot,k} \leq (\|\mathbf{E}\|_2 + d(\mathbf{S})) \bar{\mathbf{f}}_{\cdot,k}^T \bar{\mathbf{f}}_{\cdot,k}.$$

Since $\mathbf{P}_S^r(G) \bar{\mathbf{f}}_{\cdot,k}$ belongs to the subspace spanned by the smallest r eigenvectors of $\mathcal{L}_S(G)$, and $(\mathbf{I} - \mathbf{P}_S^r(G)) \bar{\mathbf{f}}_{\cdot,k}$ belongs to the subspace spanned by the remaining eigenvectors, we obtain

$$\begin{aligned} \lambda_{r+1}(\mathcal{L}_S(G)) \bar{\mathbf{f}}_{\cdot,k}^T (\mathbf{I} - \mathbf{P}_S^r(G)) \bar{\mathbf{f}}_{\cdot,k} &\leq \bar{\mathbf{f}}_{\cdot,k}^T (\mathbf{I} - \mathbf{P}_S^r(G)) \mathcal{L}_S(G) (\mathbf{I} - \mathbf{P}_S^r(G)) \bar{\mathbf{f}}_{\cdot,k} \\ &= \bar{\mathbf{f}}_{\cdot,k}^T \mathcal{L}_S(G) \bar{\mathbf{f}}_{\cdot,k} - (\mathbf{P}_S^r(G) \bar{\mathbf{f}}_{\cdot,k})^T \mathcal{L}_S(G) (\mathbf{P}_S^r(G) \bar{\mathbf{f}}_{\cdot,k}) \\ &\leq \bar{\mathbf{f}}_{\cdot,k}^T \mathcal{L}_S(G) \bar{\mathbf{f}}_{\cdot,k} \\ &\leq [\|\mathcal{L}_S(G) - \mathcal{L}_S(G')\|_2 + d(\mathbf{S})] \|\bar{\mathbf{f}}_{\cdot,k}\|_2^2. \end{aligned}$$

The result follows from the observation that $\bar{\mathbf{f}}_{\cdot,k}^T (\mathbf{I} - \mathbf{P}_S^r(G)) \bar{\mathbf{f}}_{\cdot,k} = \|\bar{\mathbf{f}}_{\cdot,k} - \mathbf{P}_S^r(G) \bar{\mathbf{f}}_{\cdot,k}\|_2^2$. \blacksquare

In Theorem 10, normalization plays a direct role because \mathbf{S} affects $\delta_r(\mathbf{S})$. Note that the eigenvalue inequality implies that $\delta_r(\mathbf{S})$ can be bounded by $1/\min_{\ell} \lambda_2(\mathcal{L}_S(G_{\ell}))$. If $\mathcal{L}_S(G) \approx \mathcal{L}_S(G')$, then $d(\mathbf{S}) \approx 0$, which means that $\mathbf{S}_j \approx \mathbf{S}_{j'}$ if j and j' belongs to the same V_{ℓ} . Moreover, $\lambda_2(\mathcal{L}_S(G_{\ell}))$ is approximately a constant; otherwise, we may reduce the largest eigenvalue $\lambda_2(G_{\ell})$ by increasing the corresponding scaling factor, which reduces $\|\mathcal{L}_S(G) - \mathcal{L}_S(G')\|_2$, thus reduces $\delta_r(\mathbf{S})$. Using reasoning that is analogous to Section 3.4, we can obtain similar conclusion under the condition $\mathcal{L}_S(G) \approx \mathcal{L}_S(G')$. That is, $\delta_r(\mathbf{S})$ is approximately minimized when $\mathbf{S}_j = \bar{s}_{\ell} = m_{\ell}$ for each $j \in V_{\ell}$.

Similar to Theorem 5, we can prove the following generalization bound using Theorem 10. For simplicity, we only consider a simple kernel $\mathbf{K}_0 = \mathbf{I}$, and take $p = 1$.

Theorem 11 *Let the assumptions of Theorem 10 hold. Consider the least squares loss $\phi(\mathbf{f}_j, y_j) = \sum_{k=1}^K (f_{j,k} - \delta_{k,y_j})^2$ in (1) using the regularization condition (6) and $\mathbf{K}_0 = \mathbf{I}$. The generalization error with optimal λ can be bounded as:*

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_j, y_j) \leq 16\delta_r(\mathbf{S}) + 8\sqrt{r/n}.$$

Proof Using Theorem 10, it can be easily verified that

$$\frac{1}{m} \sum_{j=1}^m \phi(\mathbf{P}_{\mathbf{S}}^r(G) \bar{\mathbf{f}}_{j,y_j}) + \lambda \sum_{k=1}^K (\mathbf{P}_{\mathbf{S}}^r(G) \bar{\mathbf{f}}_{\cdot,k})^T \mathbf{K}^{-1} (\mathbf{P}_{\mathbf{S}}^r(G) \bar{\mathbf{f}}_{\cdot,k}) \leq \delta_r(\mathbf{S}) + \lambda m.$$

Since regularizing with $\mathbf{K}_0 = \mathbf{I}$ is equivalent with regularizing with $\mathbf{K}_0 = \mathbf{P}_{\mathbf{S}}^r(G)$, we can use $\text{tr}(\mathbf{K}) = r$. Now using this estimate in Theorem 1, we have

$$\mathbf{E}_{Z_n} \frac{1}{m-n} \sum_{j \in \bar{Z}_n} \mathbf{err}(\hat{\mathbf{f}}_{j,y_j}) \leq 16(\delta_r(\mathbf{S}) + \lambda m) + \frac{r}{\lambda n m}.$$

Optimizing over λ gives the desired bound. ■

Similar to Theorem 7, it is possible to prove a bound for general p in Theorem 11, but the estimation of $\text{tr}_p(\mathbf{K})$ is more complicated than that of $\text{tr}(\mathbf{K})$. We skip the derivation because the extra complication is not important for the purpose of this paper. Compared to Theorem 7, the advantage of dimension reduction in Theorem 11 is that the quantity $\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)$ is replaced by $\|\mathcal{L}_{\mathbf{S}}(G) - \mathcal{L}_{\mathbf{S}}(G')\|_2$, which is typically much smaller. Instead of a rigorous analysis, we shall just give a brief intuition. For simplicity we take $\mathbf{S} = \mathbf{I}$ so that we can ignore the variations caused by \mathbf{S} . The 2-norm of the symmetric error matrix $\mathcal{L}_{\mathbf{S}}(G) - \mathcal{L}_{\mathbf{S}}(G')$ is its largest eigenvalue, which is no more than the largest 1-norm of one of its row vectors. In contrast, $\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)$ behaves similar to the absolute sum of entries of the error matrix, which is m times more than the averaged 1-norm of its row vectors. Therefore if error is relatively uniform across rows, then $\mathbf{cut}(\mathcal{L}_{\mathbf{S}}, y)$ can be at an order of m times more than $\|\mathcal{L}_{\mathbf{S}}(G) - \mathcal{L}_{\mathbf{S}}(G')\|_2$.

References

- Rie K. Ando and Tong Zhang. Learning on graph with Laplacian regularization. In *NIPS 19*, 2007.
- Mikhail Belkin and Partha Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, Special Issue on Clustering:209–239, 2004.
- Fan R.K. Chung. *Spectral Graph Theory*. Regional Conference Series in Mathematics. American Mathematical Society, Rhode Island, 1998.
- Chris Ding, Xiaofeng He, Hongyuan Zha, Ming Gu, and Horst Simon. A min-max cut algorithm for graph partitioning and data clustering. In *IEEE Int'l Conf. Data Mining*, pages 107–114, 2001.
- Daniel B. Graham and Nigel M. Allinson. Characterizing virtual eigensignatures for general purpose face recognition. *Face Recognition: From Theory to Applications, NATO ASI Series F, Computer and Systems Sciences*, 163:446–456, 1998.
- Lars Hagen and Andrew B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Transactions on Computer-Aided Design*, 11(9):1074–1085, 1992.
- Gert R.G. Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I. Jordan. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.

- Marina Meilă, Susan Shortreed, and Liang Xu. Regularized spectral learning. In *AISTATS*, 2005.
- Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *NIPS 14*, pages 849–856, 2002.
- Ralph Tyrrell Rockafellar. *Convex analysis*. Princeton University Press, Princeton, NJ, 1970.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22:888–905, 2000.
- Martin Szummer and Tommi Jaakkola. Partially labeled classification with Markov random walks. In *Advances in Neural Information Processing Systems 14*, 2002.
- Ulrike von Luxburg, Olivier Bousquet, and Mikhail Belkin. Limits of spectral clustering. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 857–864. MIT Press, Cambridge, MA, 2005.
- Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart Russell. Distance metric learning, with application to clustering with side-information. In *NIPS 15*, 2003.
- Tong Zhang. Leave-one-out bounds for kernel methods. *Neural Computation*, 15:1397–1437, 2003.
- Tong Zhang and Rie K. Ando. Analysis of spectral kernel design based semi-supervised learning. In *NIPS 18*, 2006.
- Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schlkopf. Learning with local and global consistency. In *Advances in Neural Information Processing Systems 16*, pages 321–328, 2004.
- Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *ICML 2003*, 2003.