

Integration of Multiple Networks for Robust Label Propagation

Tsuyoshi Kato*, Hisashi Kashima† and Masashi Sugiyama‡

Abstract

Transductive inference on graphs such as label propagation algorithms is receiving a lot of attention. In this paper, we address a label propagation problem on multiple networks and present a new algorithm that automatically integrates structure information brought in by multiple networks. The proposed method is robust in that irrelevant networks are automatically deemphasized, which is an advantage over Tsuda et al.’s approach [14]. We also show that the proposed algorithm can be interpreted as an EM algorithm with a Student- t prior. Finally, we demonstrate the usefulness of our method in protein function prediction.

Keywords

Label Propagation, Multiple Networks, EM Algorithm, Student- t Distribution, Protein Function Classification

1 Introduction

Recently, network-structured data is becoming increasingly popular in the field of machine learning [2, 20, 14, 12, 19, 3, 7]. Network-structured data is usually represented as an undirected graph, each of whose nodes represents an example, and each of whose edges represents a relationship between two examples. For example, in protein networks, proteins are represented as nodes, and relationships among proteins such as physical interactions and expression similarities are represented as edges. If we want to predict the functions of the proteins in a network, the task is essentially formalized as a classification problem on a network. Since we are usually given both labeled examples and unlabeled examples prior to the training stage, the task can be handled as a semi-supervised problem or a transduction problem ([17, 2, 6, 20, 18, 14]). Learning in such a setting is called *graph-based learning* (Section 2).

A generally accepted approach to graph-based learning is *label propagation* [20] (Section 3). Label propagation assumes that a pair of nodes connected by an edge should have similar predictions, and the resultant optimization problem is easily solved in a closed form.

Let us consider the situations where we are given multiple data sources, or multiple networks. In the case of protein function prediction, a wide variety of data sources are available, such as gene expression data, amino acid sequences, phylogenetic profiles, and subcellular locations. For label propagation, we have multiple networks corresponding to those multiple data sources. Since different data sources are likely to contain different information, we expect that effective integration of the complementary pieces of information will enhance the predictive performance.

To integrate multiple networks, a natural choice is to take a weighted sum of the graph Laplacians [19]. Tsuda et al. [14] proposed a method which determines the weights automatically. However, as expected from the characteristics of support vector machines (SVMs) [11], their algorithm assigns large weights to networks relevant to the task. Therefore, Tsuda et al.’s method is not robust in a situation where noisy networks are included in the set of given networks (See also Section 6).

In Section 4, we present a new robust transductive learning method that makes predictions by integrating different networks. Similar to Tsuda et al.’s method, our algorithm optimizes the weights in a linear combination of graph Laplacians. The advantage of our algorithm is that the weights are chosen so that informative data sources for prediction are automatically emphasized even in the presence of irrelevant networks. Section 5 justifies our algorithm by using a probabilistic framework for emphasizing the weights of relevant networks. Our probabilistic model employs the Student- t distribution [9], which provides robust predictions. The probabilistic model based on the Student- t distribution can be interpreted as a latent variable model in which the expectations of the latent variables are naturally derived as the weights of the graph Laplacians. Accordingly, our algorithm automatically yields an integrated network with a statistically estimated linear combination. In Section 7, we conduct experiments on label propagation for protein function classification. The experimental results reveal that the proposed methods are promising. The final section summarizes this study and discusses an extension and future work.

*GSFS, University of Tokyo

†Tokyo Research Laboratory, IBM Research

‡Tokyo Institute of Technology

Notation In this paper, vectors are denoted by bold-faced lower-case letters and matrices by bold-faced upper-case letters. Elements of vectors and matrices are not bold-faced. The transposition of a matrix \mathbf{A} is denoted by \mathbf{A}^\top , and the inverse of \mathbf{A} is by \mathbf{A}^{-1} . The $n \times n$ identity matrix is denoted by \mathbf{I}_n . The n -dimensional vector all of whose elements are one is denoted by $\mathbf{1}_n$. We use \mathbb{R} to denote the set of real numbers, \mathbb{R}^n to denote the set of n -dimensional real vectors, and $\mathbb{R}^{m \times n}$ to denote the set of $m \times n$ real matrices. The set of real nonnegative numbers is denoted by \mathbb{R}_+ , and the set of n -dimensional real nonnegative vectors is by \mathbb{R}_+^n . We use \mathbb{S}^n to denote the set of symmetric $n \times n$ matrices, \mathbb{S}_+^n to denote the set of symmetric positive semi-definite $n \times n$ matrices, and \mathbb{S}_{++}^n to denote the set of symmetric strictly positive definite $n \times n$ matrices. The symbols \leq and \geq are used to denote not only the standard inequalities between scalars, but also the componentwise inequalities between vectors.

2 Graph-based Learning Problem

Let us start with the definition of the graph-based learning problem, including two problem settings, *learning from a single network* and *learning from multiple networks*. Suppose that we are given n examples: the first $\ell < n$ examples are labeled by y_1, \dots, y_ℓ , where $y_i \in \{\pm 1\}$. For *learning from a single network*, we use an undirected network over n nodes. Each node represents an example. The network is described by a symmetric adjacency matrix $\mathbf{A} \in \mathbb{S}^n$. Every edge has a positive weight and no edge is a self-loop (i.e. $A_{ii} = 0$ for $\forall i$). The edge set is expressed as

$$\mathcal{E} \equiv \{(i, j) \mid A_{ij} > 0, i, j = 1, \dots, n\}.$$

In this paper, we wish to perform *learning from multiple networks*. If the number of networks is K , we have K adjacency matrices corresponding to the K different networks, denoted by \mathbf{A}_k ($k = 1, \dots, K$).

Graph-based learning determines the score vector $\mathbf{f} \in \mathbb{R}^n$ from the given labels and the link information (\mathbf{A} in the case of a single network, and $\mathbf{A}_1, \dots, \mathbf{A}_K$ in the case of multiple networks). For calculating the score vector $\mathbf{f} \in \mathbb{R}^n$, a typical formulation is the regularized least squares problem.

Networks are usually constructed in two stages [2]. First, we compute the distances among all the pairs of examples, and then determine the edges based on the distances. From the distances, the edge set of the network is determined by finding the k -nearest neighbors or by picking pairs with distances smaller than a threshold.

3 Existing Label Propagation Algorithm with Single Network

In this section, we review a graph-based learning approach [20] for a single network. The task is here to determine the score vector $\mathbf{f} \in \mathbb{R}^n$ from the link information \mathbf{A} and the labels \mathbf{y} . For calculating the score vector $\mathbf{f} \in \mathbb{R}^n$, the typical algorithm we consider here solves the regularized least squares problem defined by

$$(3.1) \quad \hat{\mathbf{f}} \equiv \underset{\mathbf{f}}{\operatorname{argmin}} \left(\beta_y \sum_{i=1}^{\ell} (y_i - f_i)^2 + \beta_{\text{bias}} \sum_{i=1}^n f_i^2 + \beta_{\text{net}} \sum_{(i,j) \in \mathcal{E}} A_{ij} (f_i - f_j)^2 \right)$$

where β_y , β_{bias} and β_{net} are constant. If the score is greater than a threshold, then it is classified as positive, and otherwise it is classified as negative. Using the graph Laplacian defined by $\mathbf{L} \equiv \operatorname{diag}(\mathbf{A}\mathbf{1}_n) - \mathbf{A}$, the minimization problem can be re-written as

$$(3.2) \quad \min_{\mathbf{f}} \beta_y (\mathbf{f} - \mathbf{y})^\top \mathbf{G} (\mathbf{f} - \mathbf{y}) + \mathbf{f}^\top (\beta_{\text{bias}} \mathbf{I}_n + \beta_{\text{net}} \mathbf{L}) \mathbf{f}$$

where $\mathbf{G} \in \mathbb{S}_+^n$ is a diagonal matrix with

$$(3.3) \quad G_{ii} \equiv \begin{cases} 1 & \text{if } 1 \leq i \leq \ell, \\ 0 & \text{if } \ell + 1 \leq i \leq n, \end{cases}$$

i.e. the i -th diagonal element of \mathbf{G} is one if the i -th node is labeled, or otherwise zero. Therein, we have defined $\mathbf{y} \equiv [y_1, \dots, y_\ell, 0, \dots, 0]^\top \in \mathbb{R}^n$.

4 Robust Label Propagation Algorithm

This section presents a new algorithm for robust label propagation on multiple networks. The task is to predict the score vector $\mathbf{f} \in \mathbb{R}^n$ from the labels \mathbf{y} and the network structures $\{\mathbf{A}_k\}_{k=1}^K$. A choice is integration of multiple networks followed by applying the algorithm for a single network described in the previous section. A way for integration of networks is superimposition. In order to emphasize the informative networks, we consider weighted superimposition of K networks. Let $\bar{\mathbf{u}} \in \mathbb{R}_+^K$ be the weights. The integrated adjacency matrix and the integrated graph Laplacian are then respectively given by

$$(4.4) \quad \mathbf{A}_{\text{int}}(\bar{\mathbf{u}}) = \sum_{k=1}^K \bar{u}_k \mathbf{A}_k$$

and

$$(4.5) \quad \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) = \sum_{k=1}^K \bar{u}_k \mathbf{L}_k.$$

Once the weights are determined, we obtain scores in the same way as for the standard label propagation:

$$(4.6) \quad \hat{\mathbf{f}} := \underset{\mathbf{f}}{\operatorname{argmin}} \left(\beta_y (\mathbf{f} - \mathbf{y})^\top \mathbf{G} (\mathbf{f} - \mathbf{y}) + \mathbf{f}^\top (\beta_{\text{bias}} \mathbf{I}_n + \beta_{\text{net}} \mathbf{L}_{\text{int}}(\bar{\mathbf{u}})) \mathbf{f} \right)$$

$$(4.7) \quad = \left(\mathbf{G} + \frac{\beta_{\text{bias}}}{\beta_y} \mathbf{I}_n + \frac{\beta_{\text{net}}}{\beta_y} \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) \right)^{-1} \mathbf{G} \mathbf{y}.$$

How do we compute the weights? We determine the weights iteratively with the following update rule using the current \mathbf{f} :

$$(4.8) \quad \bar{u}_k := \frac{\nu + n}{\nu + \beta_{\text{net}} \mathbf{f}^\top \mathbf{L}_k \mathbf{f}},$$

where ν is a positive constant. In Section 5.4, we will show that the update rule is naturally derived from EM algorithm. Note that

$$(4.9) \quad \mathbf{f}^\top \mathbf{L}_k \mathbf{f} = \sum_{(i,j) \in \mathcal{E}_k} A_{ij}^{(k)} (f_i - f_j)^2,$$

where $A_{ij}^{(k)}$ is the (i, j) th element in \mathbf{A}_k . If a node is likely to belong to the same class as an adjacent node, label propagation algorithms work well. In other words, an informative network should have the property that adjacent nodes tend to have similar predictions. For that reason, the value of $\mathbf{f}^\top \mathbf{L}_k \mathbf{f}$ is large for a network irrelevant to the task, whereas $\mathbf{f}^\top \mathbf{L}_k \mathbf{f}$ is small for a relevant network. Since the term $\mathbf{f}^\top \mathbf{L}_k \mathbf{f}$ is in the denominator of Eq. (4.8), the weights of relevant networks become large and the weights of irrelevant networks become small. Our algorithm is summarized in Fig. 1.

5 Probabilistic Interpretation

In this section, we give a probabilistic interpretation of our algorithm. We begin by presenting a probabilistic model for label propagation with a single network, and show that the MAP estimation coincides with the solution of the label propagation approach. We next extend the probabilistic model to the case of integration of multiple networks with fixed weights. We then introduce a prior distribution of the weights. Finally we derive an EM algorithm for MAP estimation according to that probabilistic model, and show the equivalence between the EM algorithm and the iterative algorithm described in Section 4.

5.1 Label Propagation with Single Network

We here give a probabilistic interpretation of label propagation with a single network. The label propagation

Set the initial weights:

$$\bar{\mathbf{u}} := \mathbf{1}_K / K$$

repeat

Update the scores:

$$\hat{\mathbf{f}} := \left(\mathbf{G} + \frac{\beta_{\text{bias}}}{\beta_y} \mathbf{I}_n + \frac{\beta_{\text{net}}}{\beta_y} \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) \right)^{-1} \mathbf{G} \mathbf{y}$$

Update the weights:

$$\bar{u}_k := \frac{\nu + n}{\nu + \beta_{\text{net}} \mathbf{f}^\top \mathbf{L}_k \mathbf{f}}$$

until convergence.

Figure 1: Robust Label Propagation Algorithm.

method can be viewed as performing MAP estimation of the score vector \mathbf{f} in the probabilistic model described below. The score vector $\mathbf{f} \in \mathbb{R}^n$ is in the set of model parameters. The observations y_1, \dots, y_ℓ are drawn according to the Gaussian distribution

$$(5.10) \quad p(y_i | \mathbf{f}) = \mathcal{N} \left(y_i; f_i, \frac{1}{\beta_y} \right),$$

where $\mathcal{N}(\mathbf{x}; \mathbf{m}, \mathbf{S})$ is a Gaussian density function of the observation $\mathbf{x} \in \mathbb{R}^n$ with mean $\mathbf{m} \in \mathbb{R}^n$ and covariance $\mathbf{S} \in \mathbb{S}_{++}^n$ defined by

$$(5.11) \quad \mathcal{N}(\mathbf{x}; \mathbf{m}, \mathbf{S}) \equiv \frac{1}{(2\pi)^{n/2} |\mathbf{S}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{x} - \mathbf{m})^\top \mathbf{S}^{-1} (\mathbf{x} - \mathbf{m}) \right).$$

The prior distribution of the model parameters is defined by the multivariate Gaussian distribution

$$(5.12) \quad p(\mathbf{f}) = \mathcal{N} \left(\mathbf{f}; \mathbf{0}_n, (\beta_{\text{bias}} \mathbf{I}_n + \beta_{\text{net}} \mathbf{L})^{-1} \right).$$

MAP estimation finds the value of the model parameters \mathbf{f} which maximizes the posterior probability

$$(5.13) \quad p(\mathbf{f} | \mathbf{y}) = \frac{p(\mathbf{f}) \prod_{i=1}^{\ell} p(y_i | \mathbf{f})}{p(\mathbf{y})}.$$

Since the denominator of Eq. (5.13) is constant for maximization, MAP estimation is equivalent to maximizing

the following objective function

$$(5.14) \quad \begin{aligned} & \log p(\mathbf{f}) + \sum_{i=1}^{\ell} \log p(y_i | f_i) \\ &= -\frac{1}{2} (\beta_y (\mathbf{f} - \mathbf{y})^\top \mathbf{G} (\mathbf{f} - \mathbf{y}) + \mathbf{f}^\top (\beta_{\text{bias}} \mathbf{I}_n + \beta_{\text{net}} \mathbf{L}) \mathbf{f}) \\ & \quad + \text{const.} \end{aligned}$$

We can see that the values of \mathbf{f} at the maximum of the posterior probability function are equal to the solution of Eq. (3.2). Thus the probabilistic interpretation is established.

5.2 Label Propagation with Fixed-Weight Network Integration

We next extend the probabilistic model introduced above to the integration of multiple networks with fixed weight $\bar{\mathbf{u}} \in \mathbb{R}_+^K$. A probabilistic model associated with this weighted integration is given by the conditional probabilities

$$(5.15) \quad p(y_i | \mathbf{f}) = \mathcal{N}\left(y_i; f_i, \frac{1}{\beta_y}\right), \quad i = 1, \dots, \ell$$

and the prior

$$(5.16) \quad p(\mathbf{f}) = \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, (\beta_{\text{bias}} \mathbf{I}_n + \beta_{\text{net}} \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}))^{-1}\right).$$

Using a small constant ϵ such that $\beta_{\text{bias}} > \epsilon > 0$, we add a small positive value to the diagonal elements of the graph Laplacian

$$(5.17) \quad \tilde{\mathbf{L}}_k \equiv \mathbf{L}_k + \epsilon \mathbf{I}_n.$$

Since any graph Laplacian is positive semi-definite, the matrices $\tilde{\mathbf{L}}_k$ are strictly positive definite. Exploiting this feature, we change the prior as follows:

$$(5.18) \quad p(\mathbf{f}) = \frac{1}{Z'} \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{bias}}} \mathbf{I}_n\right) \prod_{k=1}^K \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{net}} \bar{u}_k} \tilde{\mathbf{L}}_k^{-1}\right),$$

where Z' is a normalizing constant. Thus we see that the prior of the model parameters for the fixed weight integration of multiple networks is expressed as the product-of-Gaussians [1]. This rearrangement facilitates development of the probabilistic version of our robust label propagation algorithm on multiple networks.

5.3 Prior Distribution over Network Weights

We have seen the probabilistic model for label propagation with given weights. Here let us consider the situation where the weights are unknown. We introduce

a prior distribution of the weights and marginalize out the random variables of the weights from the expressions. The weights are finally obtained as the expected values. We employ the Gamma distribution for the prior of the weights. The Gamma distribution is defined by

$$(5.19) \quad \text{Gamma}(u; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} u^{\alpha-1} \exp(-\beta u)$$

for $u \geq 0, \alpha \geq 0, \beta \geq 0$. In the probabilistic model described here, each component of a network in Eq. (5.18),

$$\mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{net}} \bar{u}_k} \tilde{\mathbf{L}}_k^{-1}\right),$$

is replaced with an infinite mixture of Gaussians:

$$(5.20) \quad \int_0^\infty du_k \text{Gamma}\left(u_k; \frac{1}{2}\nu, \frac{1}{2}\nu\right) \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{net}}} (u_k \tilde{\mathbf{L}}_k)^{-1}\right)$$

where ν is a positive constant. In this model, the mixture coefficients are expressed by the Gamma distribution, and the weights $\mathbf{u} \equiv [u_1, \dots, u_K]^\top$ can be viewed as latent variables. If ν is chosen to be smaller, the prior of weights is flatter.

Our probabilistic model can be seen to employ a robust prior for the score vector. The function in Eq. (5.20) is equal to the Student- t density function [9] with zero-mean, covariance $\beta_{\text{net}}^{-1} \tilde{\mathbf{L}}_k^{-1}$, and the degree-of-freedom ν , where the density function of the Student- t distribution is generally defined by

$$(5.21) \quad \begin{aligned} & \mathcal{T}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu) \\ &= \frac{\Gamma\left(\frac{\nu+n}{2}\right) \sqrt{\det \boldsymbol{\Sigma}}}{(\sqrt{\pi\nu})^n \Gamma\left(\frac{\nu}{2}\right) \left(\sqrt{1 + \|\mathbf{x} - \boldsymbol{\mu}\|_{\boldsymbol{\Sigma}^{-1}}^2 / \nu}\right)^{\nu+n}}, \end{aligned}$$

with the mean $\boldsymbol{\mu} \in \mathbb{R}^n$, covariance $\boldsymbol{\Sigma} \in \mathbb{S}_{++}^n$, and the degree of freedom $\nu \in \mathbb{R}_+$. Thus, the prior can be viewed as the product of Student- t distributions [16] given by

$$(5.22) \quad p(\mathbf{f}) = \frac{1}{Z} \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{bias}}} \mathbf{I}_n\right) \prod_{k=1}^K \mathcal{T}\left(\mathbf{f}; \mathbf{0}_n, \beta_{\text{net}}^{-1} \tilde{\mathbf{L}}_k^{-1}, \nu\right)$$

where Z is a normalizing constant. Student- t distributions are often used for modeling noisy data robustly [4], since heavy-tailed distributions such as Student- t distributions are robust against outliers. The algorithm described below offers a robust tool for label propagation, implicitly exploiting the heavy-tailed property of Student- t distributions.

5.4 EM Algorithm for MAP Estimation We devised an EM algorithm for MAP estimation of \mathbf{f} according to the model with the prior in Eq. (5.22). The constants ν , β_y , β_{bias} and β_{net} are determined in advance. MAP estimation finds the model parameters which maximize the objective function:

$$(5.23) \quad J_{\text{lp}}(\mathbf{f}) \equiv \log p(\mathbf{f}) + \sum_{i=1}^{\ell} \log p(y_i | f_i).$$

From Eq. (5.20), the logarithm of the prior is re-written as

$$(5.24) \quad \log p(\mathbf{f}) = \log Z + \log \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{\beta_{\text{bias}}} \mathbf{I}_n\right) + \sum_{k=1}^K \log \int_0^\infty du_k h_k(\mathbf{f}, u_k),$$

where Z is a normalizing constant, and the function $h_k(\cdot, \cdot)$ is defined by

$$(5.25) \quad h_k(\mathbf{f}, u_k) \equiv \text{Gamma}\left(u_k; \frac{1}{2}\nu, \frac{1}{2}\nu\right) \mathcal{N}\left(\mathbf{f}; \mathbf{0}_n, \frac{1}{u_k \beta_{\text{net}}} \tilde{\mathbf{L}}_k^{-1}\right).$$

We introduce an arbitrary distribution $q(u_k)$ such that

$$(5.26) \quad q(u_k) \geq 0 \text{ and } \int_0^\infty du_k q(u_k) = 1.$$

Using Jensen's inequality, each component in the second term in Eq. (5.24) is bounded from below as

$$(5.27) \quad \begin{aligned} & \log \int_0^\infty du_k h_k(\mathbf{f}, u_k) \\ &= \log \int_0^\infty du_k q(u_k) \frac{h_k(\mathbf{f}, u_k)}{q(u_k)} \\ &\geq \int_0^\infty du_k q(u_k) \log \frac{h_k(\mathbf{f}, u_k)}{q(u_k)} \\ &= \int_0^\infty du_k q(u_k) \log h_k(\mathbf{f}, u_k) + \mathcal{H}[q(u_k)], \end{aligned}$$

where $\mathcal{H}[q(u_k)]$ denotes the entropy of the distribution $q(u_k)$. The inequality holds with equality when the density function $q(u_k)$ maximizes the lower-bound [10]. Specifically,

$$(5.28) \quad \begin{aligned} & \log \int_0^\infty du_k h_k(\mathbf{f}, u_k) \\ &= \int_0^\infty du_k \hat{q}(u_k) \log h_k(\mathbf{f}, u_k) + \mathcal{H}[\hat{q}(u_k)], \end{aligned}$$

where $\hat{q}(u_k)$ is the optimal distribution. Our EM algorithm attempts to maximize the lower-bound in Eq. (5.27). The EM algorithm consists in E-step and M-step: E-step computes the optimal distribution $\hat{q}(u_k)$, and M-step maximizes the logarithm of the posterior probability function with respect to the model parameters \mathbf{f} . Note that Eq. (5.28) implies that the logarithm of the posterior probability function is equal to the lower-bound in the M-step. Variational analysis derives the optimal distribution $\hat{q}(u_k)$ to be computed in the E-step as

$$(5.29) \quad \begin{aligned} & \log \hat{q}(u_k) \\ &= \log h_k(\mathbf{f}, u_k) + \text{const} \\ &= -\frac{u_k \beta_{\text{net}}}{2} \mathbf{f}^\top \tilde{\mathbf{L}}_k \mathbf{f} + \frac{n}{2} \log u_k \\ &\quad + \left(\frac{\nu}{2} - 1\right) \log u_k - \frac{\nu}{2} u_k + \text{const} \\ &= \log \text{Gamma}\left(u_k; \frac{\nu+n}{2}, \frac{\nu}{2} + \frac{\beta_{\text{net}}}{2} \mathbf{f}^\top \tilde{\mathbf{L}}_k \mathbf{f}\right), \end{aligned}$$

where 'const' denotes the terms independent of u_k . Let us denote the expectation of u_k over $\hat{q}(u_k)$ by

$$(5.30) \quad \bar{u}_k \equiv \int_0^\infty du_k \hat{q}(u_k) u_k = \frac{\nu+n}{\nu + \beta_{\text{net}} \mathbf{f}^\top \tilde{\mathbf{L}}_k \mathbf{f}}.$$

We now derive the update rule of the M-step. With the help of Eqs. (5.23), (5.24), (5.28), and (5.30), the logarithm of the posterior probability function can be re-written as

$$(5.31) \quad \begin{aligned} J_{\text{lp}}(\mathbf{f}) &= -\frac{\beta_y}{2} (\mathbf{f} - \mathbf{y})^\top \mathbf{G} (\mathbf{f} - \mathbf{y}) - \frac{\beta_{\text{bias}}}{2} \|\mathbf{f}\|^2 \\ &\quad - \frac{\beta_{\text{net}}}{2} \mathbf{f}^\top \left(\sum_{k=1}^K \bar{u}_k \tilde{\mathbf{L}}_k \right) \mathbf{f} + \text{const} \end{aligned}$$

where 'const' denotes the terms independent of \mathbf{f} . From this expression, we can see that the \bar{u}_k s play the role of the weights for the linear combination of networks. Hence \bar{u}_k s will now be called *weights*. The derivative with respect to \mathbf{f} is expressed as

$$(5.32) \quad \begin{aligned} \frac{\partial J_{\text{lp}}(\mathbf{f})}{\partial \mathbf{f}} &= -\frac{1}{2} \beta_y \mathbf{G} (\mathbf{f} - \mathbf{y}) \\ &\quad - \frac{1}{2} \left((\beta_{\text{bias}}) \mathbf{I}_n + \beta_{\text{net}} \sum_{k=1}^K \bar{u}_k \tilde{\mathbf{L}}_k \right) \mathbf{f}, \end{aligned}$$

which leads to representing the update rule for the score vector \mathbf{f} as

$$(5.33) \quad \mathbf{f} = \left(\mathbf{G} + \frac{\beta_{\text{bias}}}{\beta_y} \mathbf{I}_n + \frac{\beta_{\text{net}}}{\beta_y} \sum_{k=1}^K \bar{u}_k \tilde{\mathbf{L}}_k \right)^{-1} \mathbf{G} \mathbf{y}.$$

If we take the limit $\epsilon \rightarrow 0$, Equations (5.30) and (5.33) become Eqs (4.8) and (4.7), respectively.

The EM algorithm is then summarized as follows:

- E-step** Update \bar{u}_k using Eq. (4.8)
 for $k = 1, \dots, K$.
- M-step** Update \mathbf{f} using Eq. (4.7).

The two steps are repeated until convergence. Thus the equivalence between the iterative algorithm given in the previous section and the EM algorithm is established.

EM algorithms are guaranteed to converge to a local optimum [5]. So is our algorithm. Currently we just choose the equal weights as the initial point; we may use a multi-point search strategy to further improve the performance, although this increases the computational cost.

6 Related Work

Besides our algorithm, there exist several studies [2, 19, 14] on learning with multiple networks. Tsuda et al. [14] have proposed a label propagation algorithm. The task in their paper is exactly same as ours: to predict the labels of nodes from multiple networks. We refer to their algorithm as TSS. They pose the following optimization problem:

$$(6.34) \quad \begin{aligned} \min \quad & (\mathbf{f} - \mathbf{y})^\top (\mathbf{f} - \mathbf{y}) + c\gamma + c_0 \|\boldsymbol{\xi}\|_1 \\ \text{wrt.} \quad & \mathbf{f} \in \mathbb{R}^n, \quad \gamma \in \mathbb{R}, \quad \boldsymbol{\xi} \in \mathbb{R}_+^K \\ \text{subj. to} \quad & \mathbf{f}^\top \mathbf{L}_k \mathbf{f} \leq \gamma + \xi_k, \quad k = 1, \dots, K. \end{aligned}$$

The dual problem is given by

$$(6.35) \quad \begin{aligned} \min \quad & \mathbf{y}^\top \left(\mathbf{I}_n + \sum_k \bar{u}_k \mathbf{L}_k \right)^{-1} \mathbf{y} \\ \text{wrt.} \quad & \bar{\mathbf{u}} \in \mathbb{R}_+^K \\ \text{subj. to} \quad & \bar{\mathbf{u}} \leq c_0 \mathbf{1}_K, \quad \bar{\mathbf{u}}^\top \mathbf{1}_K \leq c. \end{aligned}$$

where \bar{u}_k is the dual variable corresponding to the inequality $\mathbf{f}^\top \mathbf{L}_k \mathbf{f} \leq \gamma + \xi_k$. Once the optimal values of the dual variables are obtained, the scores are recovered by minimizing

$$(6.36) \quad (\mathbf{f} - \mathbf{y})^\top (\mathbf{f} - \mathbf{y}) + \mathbf{f}^\top \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) \mathbf{f}$$

with respect to the score vector \mathbf{f} where

$$(6.37) \quad \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) = \sum_{k=1}^K \bar{u}_k \mathbf{L}_k.$$

Equation (6.36) implies that TSS also integrates networks and performs label propagation. The dual variables are the weights of integration. In terms of this

point, TSS is similar to our label propagation algorithm. However, TSS suffers from a serious weakness. If the weights are positive, the equality of the constraint holds. When the value of the quadratic form $\mathbf{f}^\top \mathbf{L}_k \mathbf{f}$ is less than γ , the weights \bar{u}_k are zero. As the value $\mathbf{f}^\top \mathbf{L}_k \mathbf{f}$ becomes smaller, the corresponding network is more relevant to the prediction task. For those reasons, irrelevant networks are likely to have positive weights and the weights of relevant networks vanish. Hence, TSS cannot be expected to achieve good performance when noisy networks are included.

Argyriou et al. [2] have proposed a kernel-based approach for combining different networks. Their method is based on SDP/SVM [8], which is formulated as a semidefinite programming problem, and which performs training of the SVM and optimization of the weights of the kernel matrices simultaneously. First, they convert each graph Laplacian into the Laplacian kernel matrix, and then efficiently solve the optimization problem by exploiting the sparsity of the graph Laplacians. However, their method yields an integrated kernel matrix, not the integrated graph Laplacians. Hence, if one wanted to obtain the predictions as well as the integrated network, their approach could not meet that requirement directly.

Let us discuss the computational complexity of the proposed algorithm. A Laplacian matrix is typically sparse and the number of non-zero elements is $|\mathcal{E}| + n = O(|\mathcal{E}|)$. According to Spielman and Teng's work [13], our M-step in Eq. (4.7) can be computed in $O(|\mathcal{E}|)$, while the E-step in Eq. (4.8) takes in time nearly linear in $|\mathcal{E}|$. Thus the total complexity is nearly linear in $|\mathcal{E}|R$ where R is the number of iterations. This is the same order of Tsuda et al's method (TSS) [14]. Thus our method is as computationally efficient as TSS, while ours is more robust against noisy networks.

7 Experiments

We demonstrate the utility of our label propagation algorithm with a synthetic dataset. We then show the performance of our algorithm with real biological data.

7.1 A Synthetic Example The three artificial networks we used for the demonstration are depicted in Fig. 2(a),(b),(c). Note that they have ten common nodes. We assume that the first five nodes are the positive nodes, and the last five nodes are in the negative class. Nodes 1, 3, 6 and 8 are assumed to be labeled, and the remaining nodes are unlabeled. Most of edges in \mathbf{A}_1 and \mathbf{A}_2 connect nodes with the same class labels, whereas the edges in \mathbf{A}_3 connect nodes with opposite class labels. Hence, \mathbf{A}_1 and \mathbf{A}_2 are relevant to the task, but \mathbf{A}_3 is irrelevant.

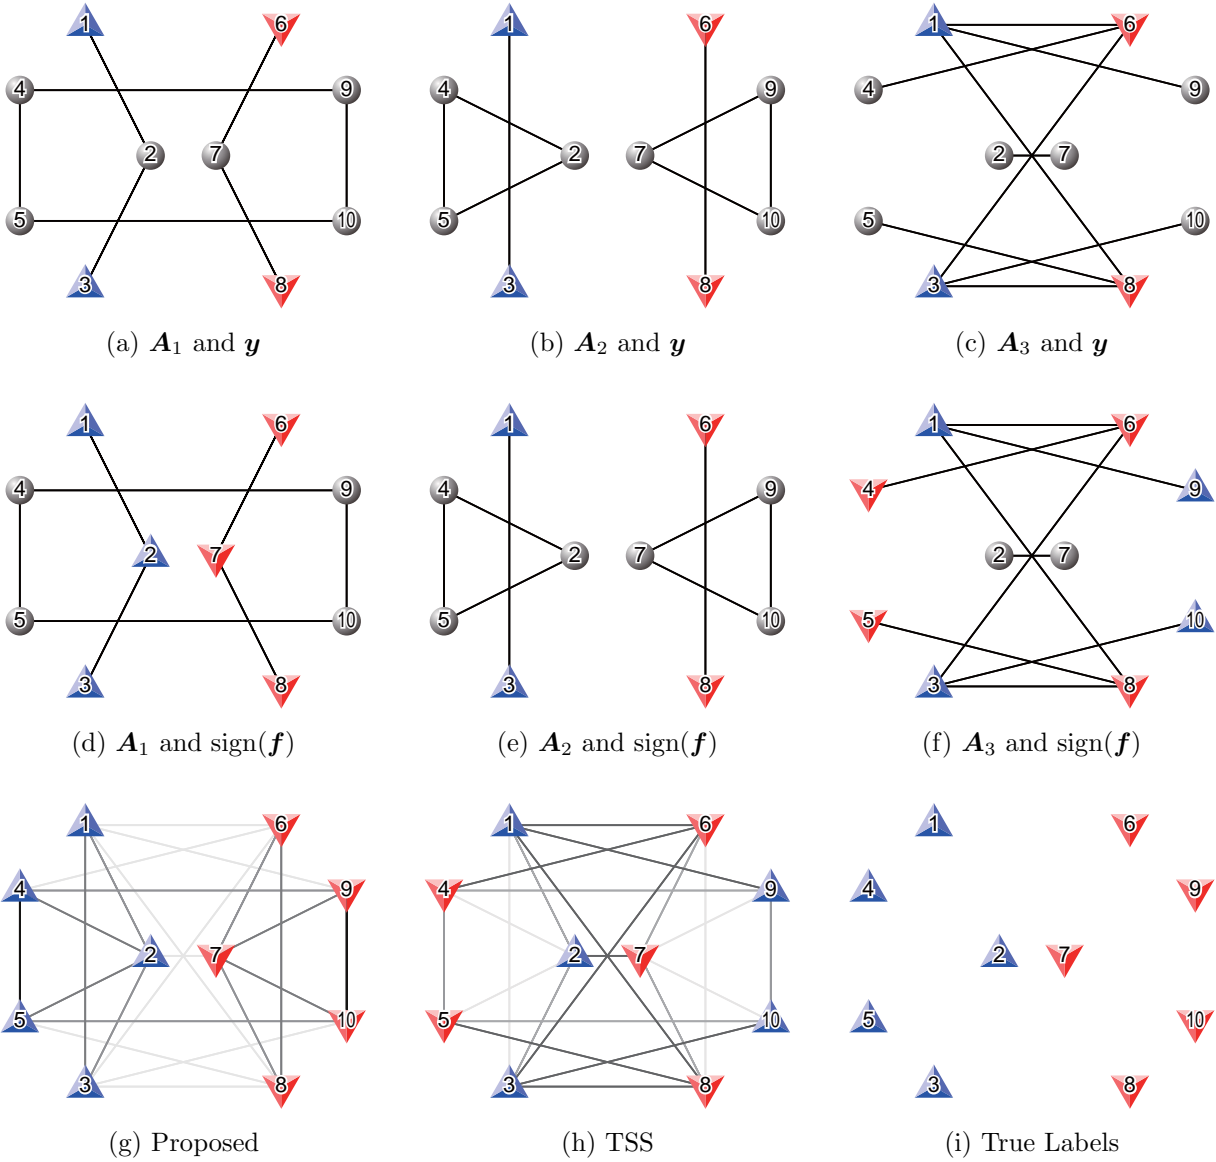


Figure 2: Here we are given three networks \mathbf{A}_1 , \mathbf{A}_2 and \mathbf{A}_3 , as depicted in (a),(b),(c) where up-pointing triangles, down-pointing triangles, and circles denote nodes with $y_i = +1$, $y_i = -1$, and unknown label, respectively. The true class labels of nodes are depicted in (i). (d),(e),(f) are the prediction results of label propagation on a individual network. The results of our algorithm and TSS are shown in (g),(h). The shade of the edges represents the network weights.

We begin with a demonstration of label propagation individually on a single network. The prediction results are shown in Fig. 2(d),(e),(f). For \mathbf{A}_1 , nodes 2 and 7 are correctly predicted, but nodes 4, 5, 9, 10 can not be predicted because these four have no paths to any labeled nodes. For \mathbf{A}_2 , no unlabeled nodes can be classified because of the same reason. For \mathbf{A}_3 , nodes 4, 5, 9, 10 are classified, but the predictions are wrong due to the irrelevant edges. The results suggest that predictions can not be made individually from any single network. We have to at least combine these networks at least for better predictions.

Now let us see the results of label propagation on the integration of networks. Figure. 2(g) depicts the prediction results of our algorithm. All the nodes are correctly classified. The resultant weights are $\bar{\mathbf{u}} = [2.98, 3.78, 0.53]^\top$. The shading of the edges in the figures represents the weights \bar{u}_k . The weight of the third network \bar{u}_3 is automatically determined to be small. As a result, the labels of all of the nodes are correctly predicted. Figure. 2(h) is the results of TSS. The values of weights produced by TSS were $\bar{\mathbf{u}} = [4.00, 2.00, 4.00]^\top$. The algorithm thus assigned a large weight to \mathbf{A}_3 , which caused poor predictions.

7.2 Protein Function Prediction Next we performed an experiment on protein function prediction [8]. (The data set is available from <http://noble.gs.washington.edu/proj/sdp-svm/>.) The task is a binary classification problem to predict whether or not each protein is ribosomal. Out of 760 proteins in total, 92 proteins are ribosomal. We used two types of input data consisting of the protein interaction network [15] and the gene expression data [8]. The gene expression data was converted to a 5-nearest-neighbor network. The two networks are denoted by \mathbf{A}_{vm} and \mathbf{A}_e , respectively. In order to test the robustness to the presence of noise, we added two decoy networks. Decoy network \mathbf{A}_{r1} was made from \mathbf{A}_{vm} by randomly shuffling the node indices. Decoy network \mathbf{A}_{r2} was made in the same way from \mathbf{A}_e . We tested two cases: in the first case only the two relevant networks \mathbf{A}_{vm} and \mathbf{A}_e were used for the primary input, and, in the second case all four networks were used. All the graph Laplacians derived from these networks were normalized so that the diagonal elements were one, and divided by the number of edges. We randomly choose 20% of the nodes as unlabeled nodes. We performed 5-fold cross validation over the labeled nodes to choose β_y , β_{bias} , β_{net} , and ν for the proposed algorithm. For performance evaluation, we performed ROC analysis while changing the threshold of classification, and computed the ROC scores (i.e. the areas under the

ROC curves).

We compared our method with two existing methods, Tsuda et al.’s method (TSS) [14], and Argyriou et al.’s kernel-based method (SDP/SVM) [2]. The hyper-parameters of TSS and SDP/SVM are also chosen using five-fold cross-validation. Table 1 summarizes the results evaluated with ROC scores averaged over ten trials. Each row represents the predictive performance as an ROC score and the weights for each of the networks. The weights for networks are denoted by \bar{u}_{vm} , \bar{u}_e , \bar{u}_{r1} , and \bar{u}_{r2} .

We used the Wilcoxon test for the statistical significance of the difference among the ROC scores. Without decoy networks, no significant differences can be observed among the three methods. Next we statistically tested whether or not each algorithm lost significant performance due to the presence of noisy networks. Adding the decoy networks, the performance of TSS degrades significantly (P-value=0.006, respectively). In particular, TSS assigns large weights to noisy networks, and thereby damages the prediction performance severely. In contrast, no significant performance loss was seen for Proposed and SDP/SVM (P-value=0.100 and 0.181, respectively).

In order to further investigate the robustness of the proposed method, we increased the number of decoy networks. Figure 3 shows the ROC scores of Proposed, TSS, and SDP/SVM methods. TSS performs considerably less well as the number of decoy networks increases. However, Proposed and SDP/SVM successfully maintain their performance.

In summary, the Proposed algorithm achieves comparable classification accuracy to SDP/SVM, yet provides the optimally integrated network $\mathbf{L}_{int}(\bar{\mathbf{u}})$.

8 Conclusion

This paper proposes using Student- t distributions to perform robust label propagation with different networks including the presence of irrelevant networks. Our algorithm consists of two phases: weighted network integration and label propagation. The network integration is done so that the weights reflect how relevant each network is to the task. The label propagation is performed on the integrated network. We described the two operations as the E-step and M-step operations, respectively, of the EM algorithm. Thus our algorithm is intuitively understandable and statistically well supported.

Although our algorithm already has the compelling property of functioning with the integrated networks, it are considering how to modify our algorithm so as to produce a sparse solution. That will be future work for this area.

Table 1: Protein function prediction using graph-learning algorithms. The upper table reports the ROC scores (areas under the ROC curves) and the weights using \mathbf{A}_{vm} and \mathbf{A}_e . The bottom table is with the two decoy networks, \mathbf{A}_{r1} , \mathbf{A}_{r2} added. The shown weights are normalized so that the sum is one.

Method	ROC score	\bar{u}_{vm}	\bar{u}_e
Proposed	1.000	0.377	0.623
TSS	0.999	0.500	0.500
SDP/SVM	0.999	-	-

Method	ROC score	\bar{u}_{vm}	\bar{u}_e	\bar{u}_{r1}	\bar{u}_{r2}
Proposed	0.998	0.279	0.408	0.16	0.154
TSS	0.721	0.200	0.000	0.400	0.400
SDP/SVM	0.999	-	-	-	-




Figure 3: Accuracy of prediction against the number of decoy networks.

We finally finish this article by presenting a possible extension of our algorithm. Extending our technique, we would perform robust manifold regularization [3]. Manifold regularization is a semi-supervised framework which uses a network to regularize a general-purpose learning algorithm. We here focus on the regularized least squares. If we have multiple networks, we can perform manifold regularization by using the integrated network. The application of our technique to manifold regularization allows us to emphasize the relevant networks and learn the prediction rule simultaneously.

Suppose we have ℓ labeled examples $\{\mathbf{x}_i, y_i\}_{i=1}^{\ell}$ and $n - \ell$ unlabeled examples $\{\mathbf{x}_i\}_{i=\ell+1}^n$ where each \mathbf{x}_i is a d -dimensional feature vector. In addition, we are given K networks for the n examples. The kernel function and the kernel matrix are denoted by $k(\cdot, \cdot)$ and $\mathbf{K} \in \mathbb{S}_{++}^n$, respectively.

Manifold regularization assumes that the score function is expressed by $f(\mathbf{x}; \boldsymbol{\alpha}) = \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{x})\alpha_i$, where $\boldsymbol{\alpha} \in \mathbb{R}^n$ is the adaptive parameter. We define the score vector function as $\mathbf{f}(\boldsymbol{\alpha}) = \mathbf{K}\boldsymbol{\alpha} \in \mathbb{R}^n$. The regularized least squares algorithm finds the parameter $\boldsymbol{\alpha} \in \mathbb{R}^n$ which minimizes

$$(8.38) \quad \beta_y \sum_{i=1}^{\ell} (y_i - f(\mathbf{x}_i; \boldsymbol{\alpha}))^2 + \beta_{\text{bias}} \boldsymbol{\alpha}^\top \mathbf{K} \boldsymbol{\alpha} + \beta_{\text{net}} \mathbf{f}(\boldsymbol{\alpha})^\top \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) \mathbf{f}(\boldsymbol{\alpha}).$$

We define the prior of $\boldsymbol{\alpha}$ as

$$(8.39) \quad p(\boldsymbol{\alpha}) = \frac{1}{Z''} \mathcal{N}\left(\boldsymbol{\alpha}; \mathbf{0}_n, \frac{1}{\beta_{\text{bias}} - \epsilon} \mathbf{K}^{-1}\right) \prod_{k=1}^K \mathcal{T}\left(\boldsymbol{\alpha}; \mathbf{0}_n, \frac{1}{\beta_{\text{net}}} \left(\mathbf{K} \tilde{\mathbf{L}}_k(\bar{\mathbf{u}}) \mathbf{K}\right)^{-1}, \nu\right),$$

where $\tilde{\mathbf{L}}_k$ is re-defined by

$$(8.40) \quad \tilde{\mathbf{L}}_k \equiv \mathbf{L}_k + \frac{\epsilon}{K\beta_{\text{net}}} \mathbf{K}^{-1}.$$

Z'' is the normalizing constant of this prior. Using a similar derivation described in Section 5.4, we can obtain the EM algorithm for obtaining the MAP estimation of the parameter $\boldsymbol{\alpha}$. Taking $\epsilon \rightarrow 0$, the update rules for the E-step and M-step are given as

$$(8.41) \quad \bar{u}_k := \frac{\nu + n}{\nu + \beta_{\text{net}} \boldsymbol{\alpha}^\top \mathbf{K} \mathbf{L}_k \mathbf{K} \boldsymbol{\alpha}}$$

and

$$(8.42) \quad \boldsymbol{\alpha} := \left(\mathbf{G} \mathbf{K} + \frac{\beta_{\text{bias}}}{\beta_y} \mathbf{I} + \frac{\beta_{\text{net}}}{\beta_y} \mathbf{L}_{\text{int}}(\bar{\mathbf{u}}) \mathbf{K} \right)^{-1} \mathbf{G} \mathbf{y},$$

respectively. We can see that, when $\mathbf{K} = \mathbf{I}_n$, this robust manifold regularization is equivalent to the robust label propagation algorithm described in Section 4. Hence, this robust manifold regularization is a generalization of the robust label propagation.

In our formulation, $p(y_i | f_i)$ is modeled by Gaussian. Employing the Student- t distribution instead may possibly increase the robustness against label errors.

References

- [1] F. V. Agakov and C. K. I. Williams. Products of Gaussians and probabilistic minor component analysis. *Neural Computation*, 14(5):1169–1182, May 2002.
- [2] A. Argyriou, M. Herbster, and M. Pontil. Combining graph Laplacians for semi-supervised learning. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 67–74. MIT Press, Cambridge, MA, 2006.
- [3] M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7:2399–2434, 2006.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer Science+Business Media, LLC, New York, USA, 2006.
- [5] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Statistical Society, Series B*, 39:1–38, 1977.
- [6] A. Kapoor, Y. A. Qi, H. Ahn, and R. Picard. Hyperparameter and kernel learning for graph-based semi-supervised classification. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 627–634. MIT Press, Cambridge, MA, 2006.
- [7] R. I. Kondor and J. Lafferty. Diffusion kernels on graphs and other discrete structures. In C. Sammut and A. G. Hoffmann, editors, *Proc. 19th International Conference on Machine Learning (ICML2002)*, pages 315–322. San Francisco, Morgan Kaufmann, 2002.
- [8] G. R. G. Lanckriet, T. D. Bie, N. Cristianini, M. Jordan, and W. Noble. A statistical framework for genomic data fusion. *Bioinformatics*, 20:2626–2635, 2004.
- [9] G. J. McLachlan and T. Krishnan. *The EM algorithm and extensions*. John Wiley & Sons, 1997.
- [10] R. M. Neal and G. E. Hinton. *Learning in Graphical Models*, chapter A view of the EM algorithm that justifies incremental, sparse, and other variants, pages 355–368. Kluwer Academic Publishers, 1998.
- [11] B. Schölkopf and A. J. Smola. *Learning with kernels*. MIT Press, Cambridge, MA, 2002.
- [12] H. Shin, J. Hill, and G. Rätsch. Graph based semi-supervised learning with sharper edges. In *ECML 2006*, pages 1008–1015, Berlin, Germany, 2006. Springer-Verlag.

- [13] D. A. Spielman and S. H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 26th annual ACM symposium on Theory of computing*, pages 81–90. ACM Press, 2004.
- [14] K. Tsuda, H. Shin, and B. Schölkopf. Fast protein classification with multiple networks. *Bioinformatics*, 21(2):i59–i65, 2005.
- [15] C. von Mering, R. Krause, B. Snel, M. Cornell, S. G. Oliver, S. Fields, and P. Bork. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 417:399–403, 2002.
- [16] M. Welling, G. Hinton, and S. Osindero. Learning sparse topographic representations with products of Student- t distributions. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 1359–1366. MIT Press, Cambridge, MA, 2003.
- [17] J. Weston, C. Leslie, D. Zhou, A. Elisseeff, and W. . Noble. Semi-supervised protein classification using cluster kernels. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
- [18] T. Zhang and R. Ando. Analysis of spectral kernel design based semi-supervised learning. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 1601–1608. MIT Press, Cambridge, MA, 2006.
- [19] T. Zhang, A. Popescul, and B. Dom. Linear prediction models with graph regularization for web-page categorization. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 821–826, Philadelphia, PA, USA, 2006.
- [20] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.