

# Protein Functional Class Prediction with a Combined Graph

Hyunjung Shin<sup>1,2</sup>, Koji Tsuda<sup>1,3</sup>, and Bernhard Schölkopf<sup>1</sup>

<sup>1</sup> Department of Empirical Inference, Max-Planck-Institute for Biological Cybernetics, Spemannstr. 38, 72076, Tübingen, Germany

{shin, tsuda, bernhard}@tuebingen.mpg.de

<sup>2</sup> Department of Industrial Engineering, Seoul National University, San 56-1, Shillim-Dong, Kwanak-Gu, 151-744, Seoul, Korea

hjshin72@snu.ac.kr

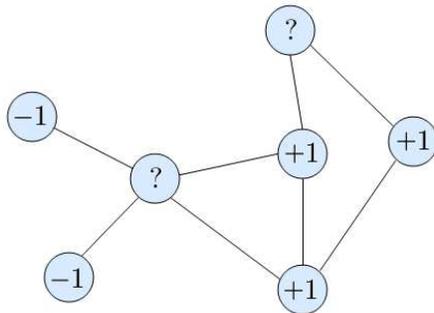
<sup>3</sup> National Institute of Advanced Industrial Science and Technology (AIST), 2-43 Aomi Koto-ku, Tokyo, Japan

**Abstract.** In bioinformatics, there exist multiple descriptions of graphs for the same set of genes or proteins. For instance, in yeast systems, graph edges can represent different relationships such as protein-protein interactions, genetic interactions, or co-participation in a protein complex, etc. Relying on similarities between nodes, each graph can be used independently for prediction of protein function. However, since different graphs contain partly independent and partly complementary information about the problem at hand, one can enhance the total information extracted by combining all graphs. In this paper, we propose a method for integrating multiple graphs within a framework of semi-supervised learning. The method alternates between minimizing the objective function with respect to network output and with respect to combining weights. We apply the method to the task of protein functional class prediction in yeast. The proposed method performs significantly better than the same algorithm trained on any single graph.

## 1 Introduction

In bioinformatics, many types of genomic data are frequently represented by using graphs of which nodes correspond to genes or proteins, and edges correspond to different relationships such as physical interactions of proteins (Schwikowski *et al.*, 2000; Uetz *et al.*, 2000; von Mering *et al.*, 2002), gene regulatory relationships (Lee *et al.*, 2002; Ihmels *et al.*, 2002; Segal *et al.*, 2003), or similarities between protein sequences (Yona *et al.*, 1999). One application using a graph representation is the prediction of protein functional class. It can be described as a binary-class classification problem on an undirected graph (see Fig.1). A protein of known class is labeled either by ‘+1’ or ‘-1’ while a protein yet unknown its class is marked as ‘?’. The goal is to predict the class of unlabeled proteins relying on similarities between nodes. Prediction of protein functional class has been studied by means of various methods such as diffusion kernel (Tsuda &

Noble, 2004a), majority vote (Hishigaki *et al.*, 2001; Schwikowski *et al.*, 2000), graph-based (Vazquez *et al.*, 2003), Bayesian (Deng *et al.*, 2003), and discriminative learning methods (Vert & Kanehisa, 2002; Lanckriet *et al.*, 2004a).



**Fig. 1.** The functional class prediction on a protein network graph: A protein of known class is labeled either by +1 or -1, and edges represent similarities between proteins. The task is to predict class of unlabeled proteins marked as ‘?’.

There can exist multiple descriptions of graphs for the same set of genes or proteins. For instance, nodes of yeast proteins can be connected in many different ways based on heterogeneous information such as protein-protein interactions, or genetic interactions, or co-participation in a protein complex, etc. Different graph sources are likely to contain partly independent and partly complementary information about the problem at hand. Thus, one can enhance the total information extracted by combining all graphs. Recently, there have arisen several methods for integrating heterogeneous data sources in bioinformatics. Most of them are based on kernel methods which represent data by means of kernel matrices defined by similarities between pairs of genes or proteins (for the kernel methods, refer to Schölkopf & Smola (2002)). Kernel matrices representing heterogeneous data types are then combined into a single matrix by various techniques. Lanckriet *et al.* (2004c) exploit semi-definite programming (SDP, see also Lanckriet *et al.* (2004b)) techniques to reduce the problem of finding optimizing kernel combinations to a convex optimization problem. This SDP-based approach yields satisfactory results when performed on genome-wide data sets, including amino acid sequences, hydrophathy profiles, gene expression data, and known protein-protein interactions. On the other hand, Kato *et al.* (2004) differentiate the worth of data sources such as ‘expensive’, which is data that is informative but difficult to obtain, and ‘cheap’, which is data that is less informative but abundantly available. Since the kernel matrix derived from the expensive data often has missing entries, they attempt to complete them using multiple cheap data. They use an expectation-maximization (EM) algorithm to simultaneously optimize the combining weights of data sources and the missing

entries of the incomplete kernel matrix (for the methodology about kernel matrix completion, refer to Tsuda *et al.* (2004b)). This EM-based method shows promising results when tested on supervised protein network inference and protein superfamily classification. The problem of multiple data sources (not only limited to graph representation) is often described as “data fusion,” which is intensely dealt with in the chapter 11–13 of the recent book of Schölkopf, Tsuda, & Vert (2004). Other methods related to integration of data sources can be found in Lanckriet *et al.* (2004a), Pavlidis *et al.* (2001), and Vert & Kanehisa (2002).

In the meantime, when data is represented as a graph, a more direct state-of-the-art in learning methods is semi-supervised learning. In semi-supervised learning, the labeled nodes provide information about the decision function, while the unlabeled nodes serve to reveal the structure of the data or data manifold by providing additional information (Chapelle *et al.*, 2003b; Zhou *et al.*, 2004a; Seeger, 2000). However, the problem of utilizing multiple data sources has yet to be explored in the framework of semi-supervised learning. In this paper, we propose a method for integrating multiple graphs within a framework of semi-supervised learning. The method alternates between minimizing the objective function with respect to network output and with respect to combining weight. We apply the method to the task of protein functional class prediction in yeast provided by the MIPS Comprehensive Yeast Genome Database (CYGD-mips.gsf.de/proj/yeast). The proposed method performs significantly better than the same algorithm trained on any single graph.

The remainder of this paper is organized as follows. In section 2, we briefly introduce semi-supervised learning and review the recent literature. Section 3 gives a detailed explanation of our proposed method. In section 4, we show experimental results. We conclude in section 5.

## 2 Semi-Supervised Learning

Let  $G = (V, E)$  denote a weighted graph where  $V = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$  is the vertex set and  $E$  is the edge set. A weight matrix associated with  $E$ , denoted as  $W$ , represents the magnitude of strength of linkage.  $W$  could be simply regarded as a non-negative similarity (or an affinity) matrix. The more similar  $\mathbf{x}_i$  to  $\mathbf{x}_j$ , the larger a value of  $w_{ij}$ . Now suppose that  $p$  vertices of  $V$  are labeled  $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_p, y_p)$  where  $y_i \in \{-1, 1\}$ , and the remaining  $q$  vertices  $\mathbf{x}_{p+1}, \mathbf{x}_{p+2}, \dots, \mathbf{x}_{p+q=n}$  are unlabeled. And accordingly, let us define  $P = \{1, 2, \dots, p\}$  for the former and  $Q = \{p+1, p+2, \dots, n\}$  for the latter. The goal of semi-supervised learning is to label those unlabeled vertices by exploiting the structure of the graph under the assumption that a label of an unlabeled vertex is more likely to be that of more adjacent or more strongly connected vertex.

To formulate the idea, let us define a function  $f : V \rightarrow \mathfrak{R}$  on  $G$  that estimates labels with this property. Then, (A) a label  $f_i$  or  $f(\mathbf{x}_i)$  estimated from  $f$  should not be too different from  $f_j$ 's of adjacent vertices (B) under the constraints  $f_i \equiv y_i, i = 1, \dots, p$ . One can obtain  $f$  by minimizing the following quadratic function

$$\sum_{i \sim j} w_{ij} (f_i - f_j)^2 + \mu \sum (f_i - y_i)^2, \quad (1)$$

where  $i \sim j$  means  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are adjacent. The first term implies the ‘smoothness’ of (A) and the second term corresponds to the ‘loss function’ of (B). Alternative functions of smoothness or loss can be found in Chapelle *et al.* (2003a). For technical convenience, a condition  $\sum f_i = 0$  can be added to Eq.(1) (Belkin & Niyogi, 2003a; Belkin *et al.*, 2003b). Very often, the quadratic problem of Eq.(1) is represented in terms of matrix,

$$\min_{\mathbf{f}} \mathbf{f}^T L \mathbf{f} + \mu (\mathbf{f} - \mathbf{y})^T (\mathbf{f} - \mathbf{y}), \quad (2)$$

where  $\mathbf{y} = [ \mathbf{y}_P^T \ \mathbf{y}_Q^T ]^T$ ,  $y_p \in \{-1, 1\}$ ,  $y_q \in \{0\}$ ,  $p \in P$ ,  $q \in Q$ , and  $\mathbf{f} = [ \mathbf{f}_P^T \ \mathbf{f}_Q^T ]^T$ ,  $f \in \mathfrak{R}$ .  $\mu$  is a parameter that trades off loss versus smoothness. The Laplacian is defined as  $L = D - W$  where  $D = \text{diag}(d_i)$ ,  $d_i = \sum_j w_{ij}$ . Instead of  $L$ , a ‘normalized Laplacian’,  $\tilde{L} = D^{-\frac{1}{2}} L D^{\frac{1}{2}}$  can be used which has many nice properties (Chung, 1997). The solution to the quadratic problem can be obtained in a form

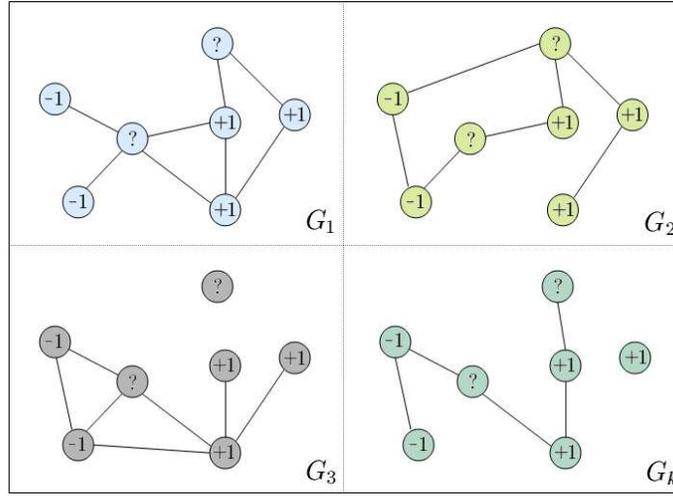
$$\mathbf{f} = \mu \{L + \mu I\}^{-1} \mathbf{y}$$

where  $I$  is an identity matrix.

There have been various semi-supervised learning algorithms, such as spectral methods and clustering (Belkin & Niyogi, 2004; Chapelle *et al.*, 2003a; Joachims, 2003; Ng *et al.*, 2001; Seeger, 2000), graph s-t mincuts (Blum & Chawla, 2001) or multi-way cuts (Kleinberg and Tardos, 1999), co-training (Blum & Mitchell, 1998), random walks (Szummer & Jaakkola, 2001; Zhou & Schölkopf, 2004b; Zhu *et al.*, 2003), and diffusion kernels (Kandola *et al.*, 2002; Kondor and Laferty, 2002; Smola & Kondor, 2003). See also ‘transductive SVM’ introduced by Vapnik (1998) which were later refined by Bennett (1999) and Joachims (1999).

### 3 Method of Combining Graphs

Given a single graph  $G$ , we can predict  $f_q$  with (2) after transforming  $G$  into a Laplacian  $L$  or a normalized Laplacian  $\tilde{L}$ . Now, consider the case where a set of graphs  $\mathbf{G} = \{G_1, G_2, \dots, G_k\}$  is given, each of which containing different aspects of the data (see Fig.2).



**Fig. 2.** Multiple graphs: Consider the case where a set of graphs  $\mathbf{G} = \{G_1, G_2, \dots, G_k\}$  is given, each of which depicts a different aspect of the data. Each graph can solely predict the label of the unlabeled nodes marked as "?", depending on its own similarity measure between nodes. However, since different graphs contain partly independent and partly complementary pieces of information about the problem at hand, one can enhance the total information extracted about the problem by combining those graphs.

To integrate multiple graphs, we consider parameterized combinations of graphs. In particular, we form the linear combination of Laplacians

$$L(\beta) = \sum_{k=1}^K \beta_k L_k, \quad (3)$$

where the weights  $\beta_k$  are constrained to be positive to assure that each Laplacian contributes to prediction of  $\mathbf{f}$ . Inserting (3) into (2), we obtain

$$\begin{aligned} \min_{\beta, \mathbf{f}} \quad & \sum_{k=1}^K \beta_k \mathbf{f}^T L_k \mathbf{f} + \mu (\mathbf{f} - \mathbf{y})^T (\mathbf{f} - \mathbf{y}), \\ \text{s.t.} \quad & \beta \geq 0, \end{aligned} \quad (4)$$

where  $\beta = [\beta_1 \ \beta_2 \ \dots \ \beta_k]^T$ . However, since each  $L_k$  is positive definite, the value  $\beta = \mathbf{0}$  is trivially optimal. This can be avoided with the additional constraint  $\beta^T \mathbf{1} = \delta$ , which yields

$$\begin{aligned} \min_{\beta, \mathbf{f}} \quad & \sum_{k=1}^K \beta_k \mathbf{f}^T L_k \mathbf{f} + \mu (\mathbf{f} - \mathbf{y})^T (\mathbf{f} - \mathbf{y}), \\ \text{s.t.} \quad & \beta \geq 0, \ \beta^T \mathbf{1} = \delta, \end{aligned} \quad (5)$$

where  $\mathbf{1} = [1 \ 1 \ \dots \ 1]^T$ . Nonetheless, the solution for  $\boldsymbol{\beta}$  may still be more sparse than desired in that case. Namely, we always get only one non-zero weight, and all the others are zero,  $\boldsymbol{\beta} = [0 \ 0 \ \dots \ \delta \ 0]^T$ . To get a reasonable set of weights, we consider to add an extra regularizer such that the term  $\mathbf{f}^\top L(\boldsymbol{\beta}) \mathbf{f}$  penalizes all the directions more equally. If the eigenvalues of  $L(\boldsymbol{\beta})$  are  $\lambda_1, \dots, \lambda_n$ , our aim is to regularize  $L(\boldsymbol{\beta})$  such that all the eigenvalues become less variant. One way to achieve this is to penalize large eigenvalues so that they are pulled toward zero. We design the regularization term as

$$-\log \det(I - L(\boldsymbol{\beta})) = -\sum_{i=1}^n \log(1 - \lambda_i).$$

Other choices might be possible, but our basic idea is to regularize the eigenvalues instead of the weights  $\beta_k$ 's. Then, the optimization problem becomes

$$\begin{aligned} \min_{\boldsymbol{\beta}, \mathbf{f}} R(\boldsymbol{\beta}, \mathbf{f}) &= \sum_{k=1}^K \beta_k \mathbf{f}^\top L_k \mathbf{f} - \log \det(I - \sum_{k=1}^K \beta_k L_k) + \mu (\mathbf{f} - \mathbf{y})^\top C (\mathbf{f} - \mathbf{y}), \\ \text{s.t. } \boldsymbol{\beta} &\geq 0, \boldsymbol{\beta}^\top \mathbf{1} = \delta, \end{aligned} \quad (6)$$

where  $\delta < 0.5$ . In the third term corresponding to the loss function, a diagonal cost matrix  $C$  is incorporated which allows different misclassification costs, i.e.,  $c_1$  for  $y_i = +1$ , and  $c_2$  for  $y_j = -1$ ,  $i, j \in P$ .

The objective function of (6) is not jointly convex, but has nice properties: by fixing  $\boldsymbol{\beta}$ , the objective function is convex with respect to  $\mathbf{f}$ , while conversely, fixing  $\mathbf{f}$  it is convex with respect to  $\boldsymbol{\beta}$ . Now, we can jointly minimize the objective function on  $\boldsymbol{\beta}$  and  $\mathbf{f}$ . We bisect the solution process similar to ‘E-step’ and ‘M-step’ of EM algorithm, and alternatively optimize both steps (Dempster *et al.* (1977); McLachlan & Krishnan (1997)). Here, we denote them instead as ‘ $\boldsymbol{\beta}$ -step’ and ‘ $\mathbf{f}$ -step,’ respectively. The algorithm is presented in Fig.3.

**Solution of [ $\mathbf{f}$ -step]** When  $\boldsymbol{\beta}$  is fixed, the solution  $\mathbf{f}$  can be obtained by

$$\frac{\partial R(\boldsymbol{\beta}, \mathbf{f})}{\partial \mathbf{f}} \Big|_{(\boldsymbol{\beta}=\boldsymbol{\beta}^i)} = \left\{ \sum_{k=1}^K \beta_k L_k + \mu C \right\} \mathbf{f} - \mu C \mathbf{y} = 0,$$

where  $C$  is a  $(n \times n)$  diagonal cost matrix. Standard linear algebra leads to a solution of the form

$$\mathbf{f} = \mu C \left\{ \sum_{k=1}^K \beta_k L_k + \mu C \right\}^{-1} \mathbf{y} \quad (7)$$

- (1) Initialize  $\beta^i$  ( $i = 0$ ) with random value under the constraints  $(\beta^i)^T \mathbf{1} = \delta$ .
- (2) [**f**-step] Given  $\beta^i$ , find  $\mathbf{f}^i$  by minimizing  $R(\beta, \mathbf{f})$  with respect to  $\mathbf{f}$ .
- (3) [ **$\beta$** -step] Given  $\mathbf{f}^i$ , find  $\beta^{i+1}$  by minimizing  $R(\beta, \mathbf{f})$  with respect to  $\beta$ .
- (4) Return  $\mathbf{f}^i$  and  $\beta^i$  if  $\left| \frac{R(\beta^{i+1}, \mathbf{f}^{i+1}) - R(\beta^i, \mathbf{f}^i)}{R(\beta^i, \mathbf{f}^i)} \right| < \epsilon$ ,  
 $i = i + 1$  and go to step-(2) otherwise.

**Fig. 3.** Algorithm: By alternating ‘ $\beta$ -step’ and ‘ $\mathbf{f}$ -step’, the optimal solution of the combining weights and the output can be found simultaneously.

**Solution of [ $\beta$ -step]** To find the solution of  $\beta$  when given  $\mathbf{f}$ , we use the gradient descent method for minimizing  $R(\beta, \mathbf{f})$  with respect to  $\beta$ . The current  $\beta^i$  is updated to  $\beta^{i+1}$  as follows:

$$\beta^{i+1} = \beta^i - \alpha^i P \mathbf{d}\beta^i \quad (8)$$

where  $\mathbf{d}\beta^i$  is the gradient vector,

$$\mathbf{d}\beta^i = \left. \frac{\partial R(\beta, \mathbf{f})}{\partial \beta} \right|_{(\mathbf{f}=\mathbf{f}^i, \beta=\beta^i)},$$

whose  $k^{th}$  element is

$$\left. \frac{\partial R(\beta, \mathbf{f})}{\partial \beta_k} \right|_{(\mathbf{f}=\mathbf{f}^i, \beta=\beta^i)} = \mathbf{f}^T L_k \mathbf{f} + \text{tr} \left[ \left( I - \sum_{j=1}^K \beta_j L_j \right)^{-1} L_k \right]. \quad (9)$$

In (9),  $\text{tr} \left[ \left( I - \sum_{j=1}^K \beta_j L_j \right)^{-1} L_k \right]$  is the derivative of  $\frac{\partial}{\partial \beta_k} \left( \log \det \left( I - \sum_k \beta_k L_k \right) \right)$  given by the following algebra. Let  $A$  be a matrix of which each element is parameterized with respect to  $t$ . The derivative of  $\frac{\partial}{\partial t} \left( \log \det A \right)$  can be drawn by

$$\begin{aligned} \frac{\partial}{\partial t} \left( \log \det A \right) &= \sum_{i,j} \frac{\partial}{\partial A_{ij}} \left( \log \det A \right) \times \frac{\partial A_{ij}}{\partial t} \\ &= \sum_{i,j} A_{ij}^{-1} \frac{\partial A_{ij}}{\partial t} \end{aligned} \quad (10)$$

where  $\sum_{i,j} \frac{\partial}{\partial A_{ij}} (\log \det A) = \frac{\partial}{\partial A} (\log \det A) = A^{-1}$ . And if  $A$  and  $B$  are symmetric, then  $\sum_{i,j} A_{ij} B_{ij} = \text{tr}[AB]$ . Thus (10) becomes

$$\sum_{i,j} A_{ij}^{-1} \frac{\partial A_{ij}}{\partial t} = \text{tr} \left[ A^{-1} \frac{\partial A}{\partial t} \right].$$

By replacing  $A$  and  $t$  with  $(I - \sum_{k=1}^K \beta_k L_k)$  and  $\beta_k$ , respectively, we find the derivative

$$\frac{\partial}{\partial \beta_k} \left( \log \det \left( I - \sum_{k=1}^K \beta_k L_k \right) \right) = \text{tr} \left[ \left( I - \sum_{j=1}^K \beta_j L_j \right)^{-1} (-L_k) \right].$$

Going back to (8), the projection matrix  $P$  is defined as

$$P = I - \frac{1}{K} \mathbf{1} \mathbf{1}^T \quad (11)$$

where  $\mathbf{1} = [1 \dots 1]^T$ . The matrix  $P$  enables the next solution of  $\beta^i$  to satisfy the constraint  $\beta^T \mathbf{1} = \delta$  so that

$$(\beta^{i+1})^T \mathbf{1} = (\beta^i + \nabla)^T \mathbf{1} = \delta$$

where  $\nabla = -\mathbf{d}\beta^i$ . Since  $(\beta^i)^T \mathbf{1} = \delta$ ,  $\nabla$  should satisfy

$$\mathbf{1}^T \nabla = 0 \quad (12)$$

which implies  $\nabla$  has to be projected onto an orthogonal space to  $\mathbf{1}^T$ . A general formula of orthogonal projection to  $A$  when  $A\nabla = \mathbf{0}$  is

$$P = I - A^T (AA^T)^{-1} A.$$

Equation (11) results from specifying  $A$  with  $\mathbf{1}^T$  in the formula. With preconditioning of  $\nabla$  with  $P$ , we now can assure (12),

$$\mathbf{1}^T (P\nabla) = \mathbf{1}^T (I - \mathbf{1}(\mathbf{1}^T \mathbf{1})^{-1} \mathbf{1}^T) \nabla = 0.$$

The  $\alpha^i$  in (8) determines the learning rate during the update. We begin with  $\alpha^i$  set to the maximum value under the condition  $\beta_k^{i+1} \geq 0, \forall k$ , and gradually reduce the magnitude as the iterations increase.

## 4 Experiments

### 4.1 Experimental Design

Our goal is to determine functional classes of yeast proteins. We used as a gold standard, the functional catalogue provided by the MIPS Comprehensive Yeast Genome Database (CYGD-mips.gsf.de/proj/yeast). The top-level categories in the functional hierarchy produce 13 classes (see table 1). A protein can belong to several functional classes. In a total of 6355 yeast proteins, however, only 3588 have class labels. The remaining yeast proteins have uncertain function and are therefore not used in evaluation. We dealt with the prediction problem as ‘one class-versus-all others’ classification tasks, one for each functional class. See Lanckriet *et al.* (2004b) for more detail.

**Table 1.** 13 CYGD functional Classes

	Classes
1	metabolism
2	energy
3	cell cycle and DNA processing
4	transcription
5	protein synthesis
6	protein fate
7	cellular transportation and transportation mechanism
8	cell rescue, defense and virulence
9	interaction with cell environment
10	cell fate
11	control of cell organization
12	transport facilitation
13	others

The input is four different types of protein interaction graphs with proteins as nodes and interactions as edges. The graphs are represented as mostly binary matrices having non-zero entry if there is interaction between the row and column proteins, 0 otherwise. The followings are the input matrices:

$W_1$  : protein-protein interactions (MIPS physical interactions),

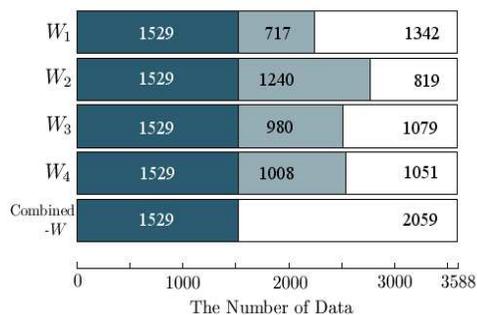
$W_2$  : genetic interactions (MIPS genetic interactions),

$W_3$  : co-participation in a protein complex (determined by tandem affinity purification, TAP), each entry is a count of the number of times two proteins appear together in a complex,

$W_4$  : co-participation in a protein complex, each entry is non-zero if and only if there is a bait-prey relationship.

There are proteins which show no interactions with others. For instance,  $W_2$  of Fig.4 has 2769 (=1529+1240) proteins with no interaction, thus only 819

(=3588-2769) are available for semi-supervised learning. And no results for the 2769 proteins remaining. Similarly, this situation arises in other graphs when they are considered individually. In contrast, using a combined graph, a protein can be used if it has at least one non-zero interaction from any graph. It amounts to the size of the union of all proteins with non-zero interaction in all graphs. In the problem at hand, 1529 of 3588 proteins have no interactions in any of the graphs. Consequently, 2059 (=3588-1529) proteins are preserved for learning. Combining graphs is also advantageous for proteins with non-zero interactions particularly when individual graphs cannot reach an accord with each other.



**Fig. 4.** The number of proteins available to learning: Dark gray indicates the number of proteins with no interaction in any of the graphs, hence unavailable to learning. Light gray indicates the number of the proteins in which there is no interaction in the specific graph but available in at least one of the other graphs. These are thus not available to learning in the specified graph but available to that of the combined graph. For each graph, the number of proteins used for learning is depicted in white.

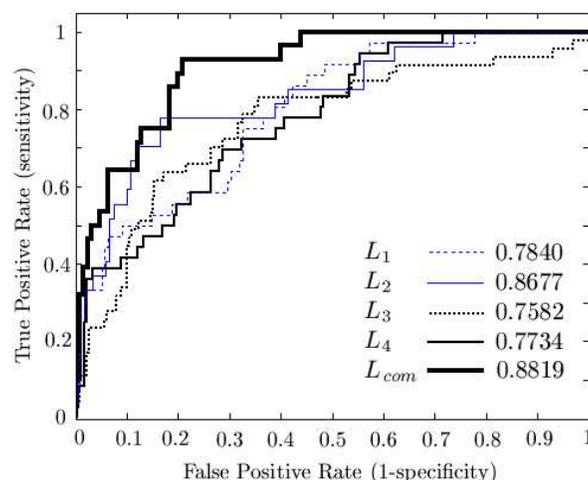
All the matrices  $W_k$  ( $k = 1, \dots, 4$ ) were transformed to ‘normalized’ Laplacian  $L_k$ ’s with dimensions of 1342, 819, 1079 and 1051, respectively. Individual Laplacians composing the combined graph, columns and rows were zero-padded up to 2059 after transformation. Hereafter, we indicate each graph with  $L_k$  ( $k = 1, \dots, 4$ ) and the combined graph with  $L_{com}$ .

The performance of  $L_{com}$  was compared with those of individual  $L_k$ ’s with the receiver operating characteristic (ROC) score, TP1FP, TP10FP, and error rate. The ROC score is the area under ROC curve (see Fig.5) that plots true positive rate (sensitivity) as a function of false positive rate (1-specificity) for differing classification thresholds (Gribskov & Robinson, 1996; Hanley & McNeil, 1982). It measures the overall quality of the ranking induced by the classifier, rather than the quality of a single value of threshold in that ranking. An ROC score of 0.5 corresponds to random guessing, and an ROC score of 1.0 implies that the algorithm succeeded in putting all of the positive examples before all of the negatives. TP1FP and TP10FP are the rates of true positives at the point

that yields 1% and 10% false positive rate on the ROC curve, respectively. Error rate is a conventional performance measurement with a fixed value threshold. Five-fold cross-validation (CV) was conducted for every class, and repeated five times in order to estimate the variance of the measurement values.

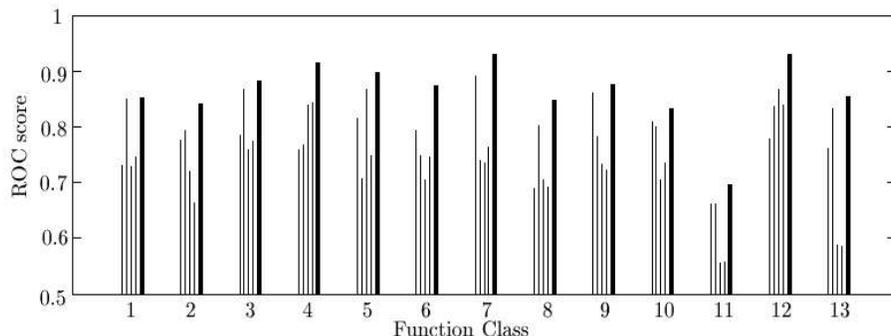
## 4.2 Results

A typical ROC curve is shown in Fig.5. The closer the curve follows the left-hand border and then the top border of the ROC space, the more accurate the classifier. The figure therefore illustrates that  $L_{com}$  is more accurate than any other single  $L_k$ . Fig.6 presents the average ROC score of each class on its test



**Fig. 5.** ROC curve: Protein functional class 3. The closer the curve follows the left-hand border and then the top border of the ROC space, the more accurate the classifier.

set when performing five-fold CV five times. The height of the stem indicates to the ROC score. Within each group of stems, a thinner stem corresponds to an individual graph in due order, such as  $L_1$ ,  $L_2$ ,  $L_3$ , and  $L_4$ , respectively while a thicker one to  $L_{com}$ . Across the 13 classes, the combined graph  $L_{com}$  outperforms any given single  $L_k$ . Overall,  $L_{com}$  yields an ROC score of 0.8313 that surpasses all those of individual  $L_k$ 's, 0.7777, 0.7836, 0.7310, and 0.7238, respectively (see Fig. 7(a)). The performances of TP1FP and TP10FP are depicted in Fig.7(b) and Fig.7(c). Among TP1FP's of  $L_k$ 's, 26.87% of  $L_2$  is the most comparable to 30.07% of  $L_{com}$ , but the gap between the best and the second best becomes larger in TP10FP by 70.15% of  $L_{com}$  and 61.22% of  $L_2$ . In Fig.7(d), the proportion of the colored bars indicates the relative weights of the different graphs when combined. Fig. 8 presents the error rates of 13 classes. A dot stands for the error rate of  $L_k$ , and the number beside it identifies the individual such as

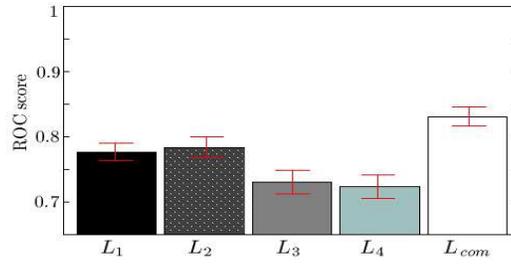


**Fig. 6.** ROC score for 13 functional protein classes: The height of the stem indicates to the ROC score. Within each group of stems, a thinner stem corresponds to an individual graph in order of such as  $L_1$ ,  $L_2$ ,  $L_3$ , and  $L_4$ , respectively while a thicker one corresponds to  $L_{com}$ . Across the 13 classes, the combined graph  $L_{com}$  outperforms any single  $L_k$ .

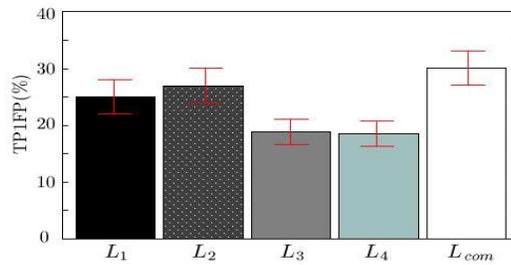
$k = 1, \dots, 4$ . The error rate of the combined graph is depicted as a square. The performance of  $L_k$  differs ‘class by class’, and the difference between the best and the worst, which is represented as a line, changes significantly as well. Therefore, it is not appropriate to put them in the order of performance. Moreover, since the difference is also large a wrong choice of graph may lead to the worst performance in specific class. On the other hand, the error rate of the combined graph is always lower than any of those of individual graphs. In addition, one does not need to take the risk involved in the choice of graphs.

To test the significance of the difference between the combined graph and individual ones, McNemar’s test was conducted (Dietterich, 1998). In principle, McNemar’s test is used to determine whether one learning algorithm outperforms another on a particular learning task. This non-parametric test could be seen as a Sign-Test in disguise. Fig.9 shows  $p$ -value distribution of McNemar’s test. The smaller  $p$ -value indicates the better the combined graph is than an individual graph, while a  $p$ -value of 1 means no statistical difference between them. A pairwise test between the combined graph and each of four graphs is conducted during five repetitions of five-fold cross-validation for 13 classes, which amounts to 1300 ( $= 4 \times 5 \times 5 \times 13$ ). For most of 1300 experiments the combined graph outperforms the individual graphs. And in 504 out of 1300 McNemar’s tests, there is a statistically significant difference between them (significance level  $\alpha=0.05$ ).

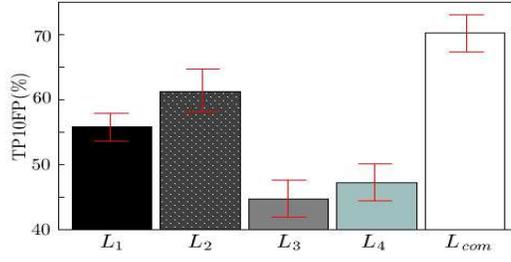
To do the comparison justice, we have only taken into consideration the proteins which are available to learning both for an individual graph and for the combined graph. For instance, when we compared  $L_{com}$  with  $L_k$ , we reported performance only on 1342 proteins (see Fig.4). However, in the combined graph, we are still able to obtain the results for another 717 proteins – that is to say,



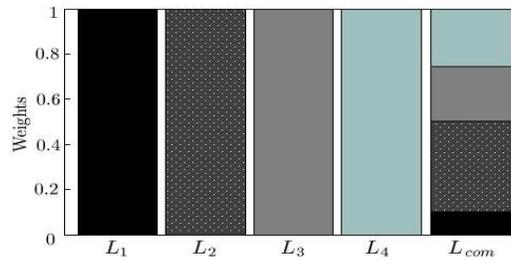
(a) ROC score



(b) TP1FP

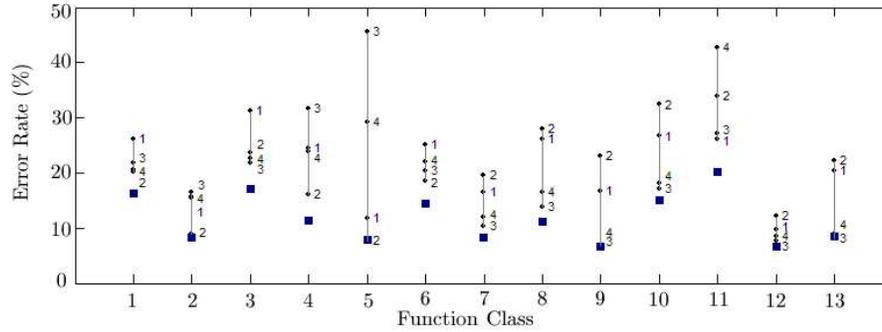


(c) TP10FP

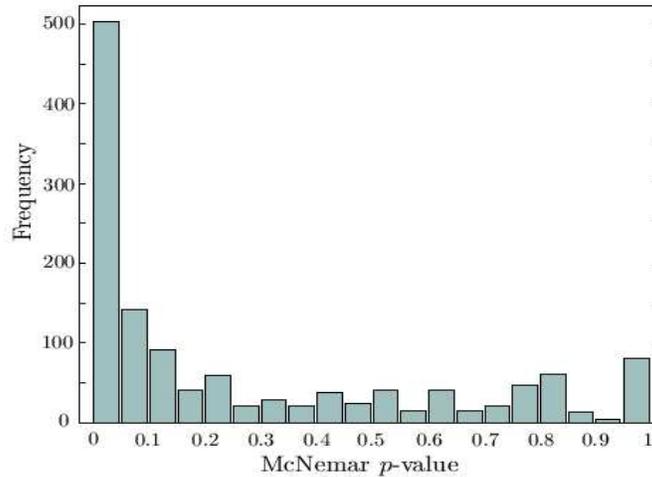


(d) Weight

**Fig. 7.** Overall performance: (a), (b), and (c) corresponds to ROC score, TP1FP and TP10FP, respectively. The height of bars indicates the average value of the measurements on five-fold CV repeated five times across 13 classes, and the error bar indicates the standard error. Seeing the results of (a), (b), and (c), the combined graph yields a better performance. In (d), the proportion of the colored bars indicates the relative weights of the different graphs when combined.



**Fig. 8.** Error rate for 13 functional protein classes: A dot stands for the error rate of  $L_k$ , and the number beside it identifies the individual,  $k = 1, \dots, 4$ . The difference between the best and the worst is represented as a line. The error rate of the combined graph is depicted as a square. The performance of  $L_k$  differs ‘class by class’, and the difference changes significantly. On the other hand, the error rate of the combined graph is always lower than any of those of individual graphs. Moreover, one does not need to take the risk involved in the choice of graphs that may lead to the worst performance in specific class.



**Fig. 9.**  $p$ -value distribution of McNemar’s test: The smaller  $p$ -value indicates a more statistically significant difference between the combined graph versus any single graph, while a  $p$ -value of 1 indicates no statistical difference between them. A pairwise test between the combined graph and each of four graphs is conducted during five repetitions of five-fold cross-validation for 13 classes, which amounts to 1300 ( $= 4 \times 5 \times 5 \times 13$ ). For most of 1300 experiments the combined graph outperforms the individual graphs. In 504 out of 1300 McNemar’s tests, there is a statistically significant difference between them (at a significance level of  $\alpha=0.05$ ).

the results of the proteins which are not available in an individual graph but available in the combined graph. Table 2 shows both error rates of the combined graph, ‘Error A’ for the former and ‘Error B’ for the latter. Error B is slightly larger than Error A, since it contains the proteins of which output is produced with fewer input graphs. Nonetheless, Error B values are still reasonable.

**Table 2.** Error rates of the combined graph: ‘Error A’ is an error rate for the proteins which are available to learning both for an individual graph and for the combined graph. ‘Error B’ contains more proteins which are not available to an individual graph but available to the combined graph. Although Error B is slightly larger than Error A, due to the relative lack of input information, it is nonetheless still a reasonable figure as an error rate.

(%)	Functional Protein Classes													Avg.
	1	2	3	4	5	6	7	8	9	10	11	12	13	
Error A	16.32	8.23	17.08	11.39	7.92	14.42	8.29	11.11	6.65	14.98	20.18	6.70	8.49	11.67
Error B	20.39	9.98	19.46	19.71	12.99	17.52	12.85	14.26	10.71	18.19	21.12	7.29	11.24	15.05

## 5 Conclusion

In this paper, we have presented a novel method for combining multiple graphs within a framework of semi-supervised learning. Similar to the EM algorithm, the method alternates between minimizing the objective function with respect to network output and with respect to combining weight. When applied to the task of functional class prediction of yeast proteins, the proposed method performed significantly better than the same algorithm trained on any single graph. The proposed method can also be used as an alternative to the model selection process. Given a single data source, it is likely to be represented in various ways by means of different parameters, i.e., different similarity measures, leading to different performances. Thus, instead of the tedious process of choosing one out of the candidate parameters, one can combine them with this method. From preliminary testing on standard data sets – Breast Cancer and Pima Indian Diabetes (from UCI Repository), we have also obtained promising results. Although the method shows good performance, it has not yet been compared with the other similar approaches such as Lanckriet *et al.* (2004a). Investigating the merits against Lanckriet *et al.* (2004a) will be the focus of future research.

## **Acknowledgement**

This work was supported by the Korean Science and Engineering Foundation (KOSEF). And the authors are grateful to Dr. Timothy Davison for his invaluable advice and proofreading.

## Bibliography

- Belkin, M. and Niyogi, P., 2003a. Using Manifold Structure for Partially Labeled Classification. *Advances in Neural Information Processing Systems (NIPS) 15*, Cambridge, MA, USA, MIT Press.
- Belkin, M., Matveeva, I., and Niyogi, P., 2003b. Regression and Regularization on Large Graphs. *University of Chicago, Computer Science, Technical Report TR-2003-11*.
- Belkin, M. and Niyogi, P., 2004. Semi-Supervised Learning on Riemannian Manifolds. *Machine Learning Journal*, to appear, (also *Technical Report TR-2001-30, Univ. of Chicago, Computer Science Dept.*)
- Bennett, K.P., 1999. Combining Support Vector and Mathematical Programming Methods for Classification, *Advances in Kernel Methods: Support Vector Learning*, B. Schölkopf et al. (Eds.), MIT press, pp. 307–326.
- Blum, A. and Mitchell, T., 1998. Combining Labeled and Unlabeled Data with Co-Training. *In Proc. of the 11th Annual Conference on Computational Learning Theory*, Morgan Kaufmann, pp. 92–100.
- Blum, A., and Chawla, S., 2001. Learning from Labeled and Unlabeled Data Using Graph Mincuts, *In Proc. of the 18th International Conference on Machine Learning (ICML)*, Morgan Kaufmann, pp. 19–26.
- Chapelle, O., Weston, J., and Schölkopf, B., 2003. Cluster Kernels for Semi-Supervised Learning. *Advances in Neural Information Processing Systems (NIPS) 15*, MIT Press, Cambridge, MA, USA, pp. 585–592.
- Chapelle, O., Schölkopf, B., and Weston, J., 2003. Semi-Supervised Learning through Principal Directions Estimation, *ICML Workshop, The Continuum from Labeled to Unlabeled Data in Machine Learning & Data Mining*
- Chung, F.R.K., (1997). *Spectral Graph Theory. Regional Conference Series in Mathematics, no. 92*.
- Dempster, A., Laird, N., and Rubin, D., 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm, *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1-38.
- Deng, M., Chen, T., and Sun, F., 2003. Integrated Probabilistic Model for Functional Prediction of Proteins, *In Proc. of the Seventh Annual International Conference on Computational Biology (RECOMB)*, ACM, pp. 95–103.
- Dietterich, T.G., 1998. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms, *Neural Computation*, vol. 10, pp. 1895–1923.
- Gribskov, M. and Robinson, N.L., 1996. Use of Receiver Operating Characteristic (ROC) Analysis to Evaluate Sequence Matching, *Computers and Chemistry*, vol. 20, no. 1, pp. 25-33.
- Hanley, J.A. and McNeil, B.J., 1982. The Meaning and Use of the Area Under a Receiver Operating Characteristic (ROC) Curve,. *Radiology*,(2nd ed), Springer, pp. 29-36.

- Hishigaki, H., Nakai, K., Ono, T., Tanigami, A., Takagi, T., 2001. Assessment of Prediction Accuracy of Protein Function from Protein-Protein Interaction Data. *Yeast*, vol. 18, no. 6, pp. 523–531.
- Ihmels, J., Friedlander, G., Bergmann, S., Sarig, O., Ziv, Y., and Barkai, N., 2002. Revealing Modular Organization in the Yeast Transcriptional Network. *Nature Genetics*, vol. 31, no. 4, pp. 370–377.
- Joachims, T., 1999. Transductive Inference for Text Classification using Support Vector Machines, *In Proc of the 16th International Conference on Machine Learning (ICML)*, Morgan Kaufmann, pp. 200–209.
- Joachims, T., 2003. Transductive Learning via Spectral Graph Partitioning, *In Proc. of the Twentieth International Conference on Machine Learning (ICML)*, AAAI press, pp. 290–297.
- Kandola, J., Shawe-Taylor, J., and Cristianini, N., 2002. Learning Semantic Similarity, *Advances in Neural Information Processing Systems (NIPS)*, 15, MIT Press, pp. 657–664.
- Kato, T., Tsuda, K., and Asai, K., 2004. Selective Integration of Multiple Biological Data for Kernel Matrix Completion, *submitted*.
- Kleinberg, J.M. and Tardos, E., 1999. Approximation Algorithms for Classification Problems with Pairwise Relationships: Metric Labeling and Markov Random Fields, *The 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 14–23.
- Kondor, R.I. and Lafferty, J., 2002. Diffusion Kernels on Graphs and Other Discrete Input Spaces, *In Proc. of the Nineteenth International Conference on Machine Learning (ICML)*, AAAI press, pp. 315–322.
- Langkriet, G.R.G., Deng, M., Cristianini, N., Jordan, M.I., and Noble, W.S., 2004. Kernel-based Data Fusion and its Application to Protein Function Prediction in Yeast, *In Pacific Symposium on Biocomputing*.
- Langkriet, G.R.G., Cristianini, N., Bartlett, El Ghaoui, L., and Jordan, M.I., 2004. Learning the Kernel Matrix with Semidefinite Programming, *Journal of Machine Learning Research*, vol 5, pp. 27–72.
- Langkriet, G.R.G., De Bie, T., Cristianini, N., Jordan, M.I., and Noble, W.S., 2004. A Statistical Framework for Genomic Data Fusion, *Bioinformatics*, vol. 20, pp. 2626–2635.
- Lee, T.I., Rinaldi, N.J., Robert, F., Odom, D.T., Bar-Joseph, Z., Gerber, G.K., Hannett, N.M., Harbison, C.T., Thompson, C.M., Simon, I., Zeitlinger, J., Jennings, E.G., Murray, H.L., Gordon, D.B., Ren, B., Wyrick, J.J., Tagne, J.B., Volkert, T.L., Fraenkel, E., Gifford, D.K., and Young, R.A., 2002. Transcriptional Regulatory Networks in *Saccharomyces cerevisiae*, *Science*, vol. 298, pp. 799–804.
- McLachlan, G. and Krishnan, T., 1997. The EM Algorithm and Extensions, *Wiley Series in Probability and Statistics*, John Wiley & Sons.
- Ng, A.Y., Jordan, M.I., and Weiss, Y., 2001. On Spectral Clustering: Analysis and an Algorithm, *Advances in Neural Information Processing Systems (NIPS)*, 14, MIT Press, pp. 849–856.
- Ngam, K., McCallum, A., Thrun, S., and Mitchell, T., 1998. Learning to Classify Text from Labeled and Unlabeled Documents, *In Proc. of the Fifteenth*

- National Conference on Artificial Intelligence (AAAI)*, AAAI press, pp. 792–799.
- Pavlidis, P., Weston, J., Cai, J., and Grundy, W.N., 2001. Gene Functional Classification from Heterogeneous Data, *In Proc of the Fifth International Conference on Computational Molecular Biology*, ACM press, pp. 242–248.
- Seeger, M., 2000. Learning with Labeled and Unlabeled Data, *Technical Report*, Edinburgh University.
- Segal, E., Shapira, M., Regev, A., Pe’er, D., Botstein, D., Koller, D., and Friedman, N., 2003. Module Networks: Identifying Regulatory Modules and Their Condition-Specific Regulators from Gene Expression Data, *Nature Genetics*, vol. 34, no. 2, pp. 166–176.
- Schölkopf, B., and Smola, A., J., 2002. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT press.
- Schölkopf, B., Tsuda, K., and Vert, J.P., 2004. *Kernel Methods in Computational Biology*, MIT Press.
- Schwikowski, B., Uetz, P., Fields, S., 2000. A Network of Protein-Protein Interactions in Yeast, *Nature Biotechnology*, vol. 18, no. 12, pp. 1257–1261.
- Smola, A.J. and Kondor, R.I., 2003. Kernels and Regularization on Graphs, *In Proc. of Conference on Learning Theory (COLT)*, Springer Verlag, pp. 144–158.
- Szummer, M. and Jaakkola, T., 2001. Partially Labeled Classification with Markov Random Walks, *Advances in Neural Information Processing Systems (NIPS)*, 14, MIT Press, pp. 945–952.
- Tsuda, K. and Noble, W.S., 2004. Learning Kernels from Biological Networks by Maximizing Entropy. *Bioinformatics*, vol. 20, pp.326–333.
- Tsuda, K., Akaho, S., and Asai, K., 2004. The *em* Algorithm for Kernel Matrix Completion with Auxiliary Data *Journal of Machine Learning Research*, vol. 4, no. 1, pp. 67–81.
- Uetz, P., Giot, L., Cagney, G., Mansfield, T.A., Judson, R.S., Knight, J.R., Lockshon, D., Narayan, V., Srinivasan, M., Pochart, P., Qureshi-Emili, A., Li, Y., Godwin, B., Conover, D., Kalbfleisch, T., Vijayadamodar, G., Yang, M., Johnston, M., Fields, S., and Rothberg, J.M., 2000. A Comprehensive Analysis of Protein-Protein Interactions in *Saccharomyces cerevisiae*, *Nature*, vol. 403, no. 6770, pp. 623–627.
- Vapnik, V., 1998. *Statistical Learning Theory*, Wiley, NY.
- Vazquez, A., Flammini, A., Maritan, A., and Vespignani, A., 2003. Global Protein Function Prediction from Protein-Protein Interaction Networks, *Nature Biotechnology*, vol. 21, no. 6, pp. 697–700.
- Vert, J.P. and Kanehisa, M., 2002. Graph-Driven Feature Extraction From Microarray Data Using Diffusion Kernels and Kernel CCA, *Advances in Neural Information Processing Systems (NIPS)* 15, Cambridge, MA, USA, MIT Press, pp. 1425–1432.
- von Mering, C., Krause, R., Snel, B., Cornell, M., Oliver, S.G., Fields, S., Bork, P., 2002. Comparative Assessment of Large-Scale Data Sets of Protein-Protein Interactions, *Nature*, vol. 23, no. 6887, pp. 399–403.

- Yona, G., Linial, N., Linial, M., 1999. ProtoMap: Automatic Classification of Protein Sequences, a Hierarchy of Protein Families, and Local Maps of the Protein Space. *Proteins*, vol. 37, no. 3, pp. 360–378.
- Zhou, D., Bousquet, O., Lal, T.N., Weston, J., and Schölkopf, B., 2004a. Learning with Local and Global Consistency, *Advances in Neural Information Processing Systems (NIPS) 16*, MIT Press, pp. 321–328.
- Zhou, D. and Schölkopf, B., 2004b. Learning from Labeled and Unlabeled Data Using Random Walks, *In Proc. of the 26th Pattern Recognition Symposium, (DAGM'04)*, also available at <http://www.kyb.tuebingen.mpg.de/bs/staff.php?gal=all>.
- Zhu, X., Ghahramani, Z., and Lafferty, J., 2003. Semi-Supervised Learning Using Gaussian Fields and Harmonic Functions, *In Proc. of the Twentieth International Conference on Machine Learning (ICML)*, AAAI press, pp. 912-919