

# UNIVERSALLY CONSISTENT VERTEX CLASSIFICATION FOR LATENT POSITIONS GRAPHS<sup>1</sup>

BY MINH TANG, DANIEL L. SUSSMAN AND CAREY E. PRIEBE

*Johns Hopkins University*

In this work we show that, using the eigen-decomposition of the adjacency matrix, we can consistently estimate feature maps for latent position graphs with positive definite link function  $\kappa$ , provided that the latent positions are i.i.d. from some distribution  $F$ . We then consider the exploitation task of vertex classification where the link function  $\kappa$  belongs to the class of universal kernels and class labels are observed for a number of vertices tending to infinity and that the remaining vertices are to be classified. We show that minimization of the empirical  $\varphi$ -risk for some convex surrogate  $\varphi$  of 0–1 loss over a class of linear classifiers with increasing complexities yields a universally consistent classifier, that is, a classification rule with error converging to Bayes optimal for any distribution  $F$ .

**1. Introduction.** The classical statistical pattern recognition setting involves

$$(X, Y), (X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \stackrel{\text{i.i.d.}}{\sim} F_{\mathcal{X}, \mathcal{Y}},$$

where the  $X_i \in \mathcal{X} \subset \mathbb{R}^d$  are observed feature vectors, and the  $Y_i \in \mathcal{Y} = \{-1, 1\}$  are observed class labels, for some probability distribution  $F_{\mathcal{X}, \mathcal{Y}}$  on  $\mathcal{X} \times \mathcal{Y}$ . Let  $\mathcal{D} = \{(X_i, Y_i)\}_{i=1}^n$ . A classifier  $h(\cdot; \mathcal{D}) : \mathcal{X} \mapsto \{-1, 1\}$  whose probability of error  $\mathbb{P}[h(X; \mathcal{D}) \neq Y | \mathcal{D}]$  approaches Bayes-optimal as  $n \rightarrow \infty$  for all distributions  $F_{\mathcal{X}, \mathcal{Y}}$  is said to be *universally consistent*. For example, the  $k$ -NN classifier with  $k \rightarrow \infty, k/n \rightarrow 0$  is universally consistent [34].

In this paper, we consider the case wherein the feature vectors are unobserved, and we observe instead a latent position graph  $G = G(X, X_1, \dots, X_n)$  on  $n + 1$  vertices with positive definite link function  $\kappa : \mathcal{X} \times \mathcal{X} \mapsto [0, 1]$ . The graph  $G$  is constructed such that there is a one-to-one relationship between the vertices of  $G$  and the feature vectors  $X, X_1, \dots, X_n$ , and the edges of  $G$  are *conditionally independent* Bernoulli random variables given the *latent*  $X, X_1, \dots, X_n$ . We show that there exists a universally consistent classification rule for this extension of the

---

Received December 2012; revised March 2013.

<sup>1</sup>Supported in part by National Security Science and Engineering Faculty Fellowship (NSSEFF), Johns Hopkins University Human Language Technology Center of Excellence (JHU HLT COE) and the XDATA program of the Defense Advanced Research Projects Agency (DARPA) administered through Air Force Research Laboratory contract FA8750-12-2-0303.

*MSC2010 subject classifications.* Primary 62H30; secondary 62C12, 62G20.

*Key words and phrases.* Classification, latent space model, convergence of eigenvectors, Bayes-risk consistency, convex cost function.

---

**Algorithm 1** Vertex classifier on graphs

---

**Input:**  $\mathbf{A} \in \{0, 1\}^{n \times n}$ , training set  $\mathcal{T} \subset [n] = \{1, 2, \dots, n\}$  and labels  $\mathbf{Y}_{\mathcal{T}} = \{Y_i : i \in \mathcal{T}\}$ .

**Output:** Class labels  $\{\hat{Y}_j : j \in [n] \setminus \mathcal{T}\}$ .

*Step 1:* Compute the eigen-decomposition of  $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{U}^T$ .

*Step 2:* Let  $d$  be the “elbow” in the scree plot of  $\mathbf{A}$ ,  $\mathbf{S}_{\mathbf{A}}$  the diagonal matrix of the top  $d$  eigenvalues of  $\mathbf{A}$  and  $\mathbf{U}_{\mathbf{A}}$  the corresponding columns of  $\mathbf{U}$ .

*Step 3:* Define  $\mathbf{Z}$  to be  $\mathbf{U}_{\mathbf{A}}\mathbf{S}_{\mathbf{A}}^{1/2}$ . Denote by  $Z_i$  the  $i$ th row of  $\mathbf{Z}$ . Define  $\mathbf{Z}_{\mathcal{T}}$  the rows of  $\mathbf{Z}$  corresponding to the indices set  $\mathcal{T}$ .  $\mathbf{Z}$  is called the adjacency spectral embedding of  $\mathbf{A}$ .

*Step 4:* Find a linear classifier  $\tilde{g}_n$  that minimizes the empirical  $\varphi$ -loss when trained on  $(\mathbf{Z}_{\mathcal{T}}, \mathbf{Y}_{\mathcal{T}})$  where  $\varphi$  is a convex loss function that is a surrogate for 0–1 loss.

*Step 5:* Apply  $\tilde{g}_n$  on the  $\{Z_j : j \in [n] \setminus \mathcal{T}\}$  to obtain the  $\{\hat{Y}_j : j \in [n] \setminus \mathcal{T}\}$ .

---

classical pattern recognition setup to latent position graph models, provided that the link function  $\kappa$  is an element of the class of *universal kernels*. In particular, we show that a classifier similar to the one described in Algorithm 1 is universally consistent. Algorithm 1 is an example of a procedure that first embeds data into some Euclidean space and then performs inference in that space. These kind of procedures are popular in analyzing graph data, as is evident from the vast literature on multidimensional scaling, manifold learning and spectral clustering.

The above setting of classification for latent position graphs, with  $\kappa$  being the inner product in  $\mathbb{R}^d$ , was previously considered in [36]. It was shown there that the eigen-decomposition of the adjacency matrix  $\mathbf{A}$  yields a consistent estimator, up to some orthogonal transformation, of the latent vectors  $X, X_1, \dots, X_n$ . Therefore, the  $k$ -NN classifier, using the estimated vectors, with  $k \rightarrow \infty, k/n \rightarrow 0$  is universally consistent. When  $\kappa$  is a general, possibly unknown link function, we cannot expect to recover the latent vectors. However, we can obtain a consistent estimator of some feature map  $\Phi : \mathcal{X} \mapsto \mathcal{H}$  of  $\kappa$ . Classifiers that use only the feature map  $\Phi$  are universally consistent if the space  $\mathcal{H}$  is isomorphic to some dense subspace of the space of measurable functions on  $\mathcal{X}$ . The notion of a universal kernel [25, 30, 31] characterizes those  $\kappa$  whose feature maps  $\Phi$  induce a dense subspace of the space of measurable functions on  $\mathcal{X}$ .

The structure of our paper is as follows. We introduce the framework of latent position graphs in Section 2. In Section 3, we show that the eigen-decomposition of the adjacency matrix  $\mathbf{A}$  yields a consistent estimator for a feature map  $\Phi : \mathcal{X} \mapsto l_2$  of  $\kappa$ . We discuss the notion of universal kernels and the problem of vertex classification using the estimates of the feature map  $\Phi$  in Section 4. In particular, we show that the classification rule obtained by minimizing a convex surrogate of the 0–1 loss over a class of linear classifiers in  $\mathbb{R}^d$  is universally consistent, provided

that  $d \rightarrow \infty$  in a specified manner. We conclude the paper with a discussion of how some of the results presented herein can be extended and other implications.

We make a brief comment on the setup of the paper. The main contribution of the paper is the derivation of the estimated feature maps and their use in constructing a universally consistent vertex classifier. We have thus considered a less general setup of compact metric spaces, linear classifiers, and convex, differentiable loss functions. It is possible to extend the results herein to a more general setup where the latent positions are elements of a (non-compact) metric space, the class of classifiers are uniformly locally-Lipschitz, and the convex loss function satisfies the classification-calibrated property [2].

**2. Framework.** Let  $(\mathcal{X}, d)$  be a compact metric space and  $F$  a probability measure on the Borel  $\sigma$ -field of  $\mathcal{X}$ . Let  $\kappa : \mathcal{X} \times \mathcal{X} \mapsto [0, 1]$  be a continuous, positive definite kernel. Let  $L^2(\mathcal{X}, F)$  be the space of square-integrable functions with respect to  $F$ . We can define an integral operator  $\mathcal{K} : L^2(\mathcal{X}, F) \mapsto L^2(\mathcal{X}, F)$  by

$$\mathcal{K} f(x) = \int_{\mathcal{X}} \kappa(x, x') f(x') F(dx').$$

$\mathcal{K}$  is a compact operator and is of trace class.

Let  $\{\lambda_j\}$  be the set of eigenvalues of  $\mathcal{K}$  ordered as  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ . Let  $\{\psi_j\}$  be a set of orthonormal eigenfunctions of  $\mathcal{K}$  corresponding to the  $\{\lambda_j\}$ , that is,

$$\mathcal{K} \psi_j = \lambda_j \psi_j,$$

$$\int_{\mathcal{X}} \psi_i(x) \psi_j(x) dF(x) = \delta_{ij}.$$

The following Mercer representation theorem [12, 33] provides a representation for  $\kappa$  in terms of the eigenvalues and eigenfunctions of  $\mathcal{K}$  defined above.

**THEOREM 2.1.** *Let  $(\mathcal{X}, d)$  be a compact metric space and  $\kappa : \mathcal{X} \times \mathcal{X} \mapsto [0, 1]$  be a continuous positive definite kernel. Let  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  be the eigenvalues of  $\mathcal{K}$  and  $\psi_1, \psi_2, \dots$  be the associated eigenvectors. Then*

$$(2.1) \quad \kappa(x, x') = \sum_{j=1}^{\infty} \lambda_j \psi_j(x) \psi_j(x').$$

The sum in equation (2.1) converges absolutely for each  $x$  and  $x'$  in  $\text{supp}(F) \times \text{supp}(F)$  and uniformly on  $\text{supp}(F) \times \text{supp}(F)$ . Let  $\mathcal{H}$  denote the reproducing kernel Hilbert space of  $\kappa$ . Then the elements  $\eta \in \mathcal{H}$  are of the form

$$(2.2) \quad \eta = \sum_j a_j \sqrt{\lambda_j} \psi_j \quad \text{with } (a_j) \in l_2,$$

and the inner product on  $\mathcal{H}$  is given by

$$(2.3) \quad \left\langle \sum_j a_j \sqrt{\lambda_j} \psi_j, \sum_j b_j \sqrt{\lambda_j} \psi_j \right\rangle_{\mathcal{H}} = \sum_j a_j b_j.$$

By Mercer’s representation theorem, we have  $\kappa(\cdot, x) = \sum_j \sqrt{\lambda_j} \psi_j(x) \times \sqrt{\lambda_j} \psi_j(\cdot)$ . We thus define the feature map  $\Phi : \mathcal{X} \mapsto l_2$  by

$$\Phi(x) = (\sqrt{\lambda_j} \psi_j(x) : j = 1, 2, \dots).$$

Let  $d$  be an integer with  $d \geq 1$ . We also define the following map  $\Phi_d : \mathcal{X} \mapsto \mathbb{R}^d$

$$\Phi_d(x) = (\sqrt{\lambda_j} \psi_j(x) : j = 1, 2, \dots, d).$$

We will refer to  $\Phi_d$  as the truncation of  $\Phi$  to  $\mathbb{R}^d$ .

Now, for a given  $n$ , let  $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} F$ . Define  $\mathbf{K} = (\kappa(X_i, X_j))_{i,j=1}^n$ . Let  $\mathbf{A}$  be a symmetric random hollow matrix where the entries  $\{\mathbf{A}_{ij}\}_{i < j}$  are *conditionally independent* Bernoulli random variables given the  $\{X_i\}_{i=1}^n$  with  $\mathbb{P}[\mathbf{A}_{ij} = 1] = \mathbf{K}_{ij}$  for all  $i, j \in [n], i < j$ .  $\mathbf{A}$  is the adjacency matrix corresponding to a graph with vertex set  $\{1, 2, \dots, n\}$ . A graph  $G$  whose adjacency matrix  $\mathbf{A}$  is constructed as above is an instance of a latent position graph [19] where the latent positions are sampled according to  $F$ , and the link function is  $\kappa$ .

2.1. *Related work.* The latent position graph model and the related latent space approach [19] is widely used in network analysis. It is a generalization of the stochastic block model (SBM) [20] and variants such as the degree-corrected SBM [21] or the mixed-membership SBM [1] or the random dot product graph model [39]. It is also closely related to the inhomogeneous random graph model [8] or the exchangeable graph model [15].

There are two main sources of randomness in latent position graphs. The first source of randomness is due to the sampling procedure, and the second source of randomness is due to the conditionally independent Bernoulli trials that gave rise to the edges of the graphs. The randomness in the sampling procedure and its effects on spectral clustering and/or kernel PCA have been widely studied. In the manifold learning literature, the latent positions are sampled from some manifold in Euclidean space and [3, 17, 18] among others studied the convergence of the various graph Laplacian matrices to their corresponding Laplace–Beltrami operators on the manifold. The authors of [28, 38] studied the convergence of the eigenvalues and eigenvectors of the graph Laplacian to the eigenvalues and eigenfunctions of the corresponding operators in the spectral clustering setting.

The matrix  $\mathbf{K}/n$  can be considered as an approximation of  $\mathcal{K}$  for large  $n$ ; that is, we expect the eigenvalues and eigenvectors of  $\mathbf{K}/n$  to converge to the eigenvalues and eigenfunctions of  $\mathcal{K}$  in some sense. This convergence is important in understanding the theoretical properties of kernel PCA; see, for example, [4, 9, 23, 28, 29, 41]. We summarized some of the results from the literature that directly pertain to the current paper in Appendix B.

The Bernoulli trials at each edge and their effects had also been studied. For example, a result in [10] on matrix estimation for noise-perturbed and subsampled

matrices showed that by thresholding the dimensions in the singular value decomposition of the adjacency matrix, one can recover an estimate  $\hat{\mathbf{K}}$  of the kernel matrix  $\mathbf{K}$  with small  $\|\hat{\mathbf{K}} - \mathbf{K}\|$ . Oliveira [26] studied the convergence of the eigenvalues and eigenvectors of the adjacency matrix  $\mathbf{A}$  to that of the integral operator  $\mathcal{K}$  for the class of inhomogeneous random graphs. The inhomogeneous random graphs in [26] have latent positions that are uniform  $[0, 1]$  random variables with the link function  $\kappa$  being arbitrary symmetric functions.

As we have mentioned in Section 1, Algorithm 1 is an example of a popular approach in multidimensional scaling, manifold learning and spectral clustering where inference on graphs proceeds by first embedding the graph into Euclidean space followed by inference on the resulting embedding. It is usually assumed that the embedding is conducive to the subsequent inference tasks. Justification can also be provided based on the theoretical results about convergence, for example, the convergence of the eigenvalues and eigenvectors to the eigenvalues and eigenfunctions of operators, or the convergence of the estimated entries, cited above. However, these justifications do not consider the subsequent inference problem; that is, these convergence results do not directly imply that inference using the embeddings are meaningful. Recently, the authors of [11, 16, 27, 35] showed that the clustering using the embeddings are meaningful, that is, consistent, for graphs based on the stochastic block model and the extended planted partition model. The main impetus for this paper is to give similar theoretical justification for the classification setting. The latent position graph model is thus a surrogate model—a widely-used model with sufficiently simple structure that allows for clear, concise theoretical results.

**3. Estimation of feature maps.** We assume the setting of Section 2. Let us denote by  $\mathcal{M}_d(\mathbb{R})$  and  $\mathcal{M}_{d,n}(\mathbb{R})$  the set of  $d \times d$  matrices and  $d \times n$  matrices on  $\mathbb{R}$ , respectively. Let  $\tilde{\mathbf{U}}_{\mathbf{A}} \tilde{\mathbf{S}}_{\mathbf{A}} \tilde{\mathbf{U}}_{\mathbf{A}}^{\top}$  be the eigen-decomposition of  $\mathbf{A}$ . For a given  $d \geq 1$ , let  $\mathbf{S}_{\mathbf{A}} \in \mathcal{M}_d(\mathbb{R})$  be the diagonal matrix comprised of the  $d$  largest eigenvalues of  $\mathbf{A}$ , and let  $\mathbf{U}_{\mathbf{A}} \in \mathcal{M}_{n,d}(\mathbb{R})$  be the matrix comprised of the corresponding eigenvectors. The matrices  $\mathbf{S}_{\mathbf{K}}$  and  $\mathbf{U}_{\mathbf{K}}$  are defined similarly. For a matrix  $\mathbf{M}$ ,  $\|\mathbf{M}\|$  refers to the spectral norm of  $\mathbf{M}$  while  $\|\mathbf{M}\|_F$  refers to the Frobenius norm of  $\mathbf{M}$ . For a vector  $v \in \mathbb{R}^n$ ,  $\|v\|$  will denote the Euclidean norm of  $v$ .

The key result of this section is the following theorem which shows that, given that there is a gap in the spectrum of  $\mathcal{K}$  at  $\lambda_d(\mathcal{K})$ , by using the eigen-decomposition of  $\mathbf{A}$  we can accurately estimate the truncated map  $\Phi_d$  in Section 2 up to an orthogonal transformation. We note that the dependence on  $F$ , the distribution of the  $\{X_i\}$ , in the following result is implicit in the definition of the spectral gap  $\delta_d$  of  $\mathcal{K} : L^2(\mathcal{X}, F) \mapsto L^2(\mathcal{X}, F)$ .

**THEOREM 3.1.** *Let  $d \geq 1$  be given. Denote by  $\delta_d$  the quantity  $\lambda_d(\mathcal{K}) - \lambda_{d+1}(\mathcal{K})$ , and suppose that  $\delta_d > 0$ . Then with probability greater than  $1 - 2\eta$ ,*

there exists a unitary matrix  $\mathbf{W} \in \mathcal{M}_d(\mathbb{R})$  such that

$$(3.1) \quad \|\mathbf{U}_A \mathbf{S}_A^{1/2} \mathbf{W} - \Phi_d\|_F \leq 27\delta_d^{-2} \sqrt{d \log(n/\eta)},$$

where  $\Phi_d$  denotes the matrix in  $\mathcal{M}_{n,d}(\mathbb{R})$  whose  $i$ th row is  $\Phi_d(X_i)$ . Let us denote by  $\hat{\Phi}_d(X_i)$  the  $i$ th row of  $\mathbf{U}_A \mathbf{S}_A^{1/2} \mathbf{W}$ . Then, for each  $i \in [n]$  and any  $\varepsilon > 0$ ,

$$(3.2) \quad \mathbb{P}[\|\hat{\Phi}_d(X_i) - \Phi_d(X_i)\| > \varepsilon] \leq 27\delta_d^{-2} \varepsilon^{-1} \sqrt{\frac{6d \log n}{n}}.$$

We now proceed to prove Theorem 3.1. A rough sketch of the argument goes as follows. First we will show that the projection of  $\mathbf{A}$  onto the subspace spanned by  $\mathbf{U}_A$  is “close” to the projection of  $\mathbf{K}$  onto the subspace spanned by  $\mathbf{U}_K$ . Then we will use results on the convergence of spectra of  $\mathbf{K}$  to the spectra of  $\mathcal{K}$  to show that the subspace spanned by  $\mathbf{U}_A$  is also “close” to the subspace spanned by  $\Phi_d$ . We note that, for conciseness and simplicity in the exposition, all probability statements involving the matrix  $\mathbf{A}$  or its related quantities, for example,  $\mathbf{U}_A, \mathbf{S}_A$ , are assumed to hold conditionally on the  $\{X_1, X_2, \dots, X_n\}$ .

We need the following bound for the perturbation  $\mathbf{A} - \mathbf{K}$  from [26]. The convergence of the spectra of  $\mathbf{A}$  to that of  $\mathcal{K}$  as given by Theorem 6.1 in [26] is similar to that given in the proof of Theorem 3.1 in the current paper, but there are sufficient differences between the two settings, and we do not see an obvious way to apply the conclusions of Theorem 6.1 in [26] to the current paper.

PROPOSITION 3.2. For  $\mathbf{A}$  and  $\mathbf{K}$  as defined above, with probability at least  $1 - \eta$ , we have

$$(3.3) \quad \|\mathbf{A} - \mathbf{K}\| \leq 2\sqrt{\Delta \log(n/\eta)} \leq 2\sqrt{n \log(n/\eta)},$$

where  $\Delta$  is the maximum vertex degree.

The constant in equation (3.3) was obtained by replacing a concentration inequality in [26] with a slightly stronger inequality from [37]. We now show that the projection matrix for the subspace spanned by  $\mathbf{U}_A$  is close to the projection matrix for the subspace spanned by  $\mathbf{U}_K$ .

PROPOSITION 3.3. Let  $\mathcal{P}_A = \mathbf{U}_A \mathbf{U}_A^T$  and  $\mathcal{P}_K = \mathbf{U}_K \mathbf{U}_K^T$ . Denote by  $\delta_d$  the quantity  $\lambda_d(\mathcal{K}) - \lambda_{d+1}(\mathcal{K})$ , and suppose that  $\delta_d > 0$ . If  $n$  is such that  $\delta_d \geq 8(1 + \sqrt{2})n^{-1/2} \sqrt{\log(n/\eta)}$ . Then with probability at least  $1 - 2\eta$ ,

$$(3.4) \quad \|\mathcal{P}_A - \mathcal{P}_K\| \leq 4\sqrt{\frac{\log(n/\eta)}{n\delta_d^2}}.$$

PROOF. By equation (B.4) in Theorem B.2, we have with probability at least  $1 - \eta$ ,

$$\frac{\lambda_d(\mathbf{K})}{n} - \frac{\lambda_{d+1}(\mathbf{K})}{n} \geq \delta_d - 4\sqrt{2}\sqrt{\frac{\log(2/\eta)}{n}}.$$

Now, let  $S_1$  and  $S_2$  be defined as

$$S_1 = \{\lambda : \lambda \geq n\lambda_d(\mathbf{K}) - 2\sqrt{n \log(n/\eta)}\},$$

$$S_2 = \{\lambda : \lambda < n\lambda_{d+1}(\mathbf{K}) + 2\sqrt{n \log(n/\eta)}\}.$$

Then we have, with probability at least  $1 - \eta$ ,

$$(3.5) \quad \begin{aligned} \text{dist}(S_1, S_2) &\geq n\delta_d - 4\sqrt{2}\sqrt{n \log(2/\eta)} - 4\sqrt{n \log(n/\eta)} \\ &\geq n\delta_d - 4(1 + \sqrt{2})\sqrt{n \log(n/\eta)}. \end{aligned}$$

Suppose for the moment that  $S_1$  and  $S_2$  are disjoint, that is, that  $\text{dist}(S_1, S_2) > 0$ . Let  $\mathcal{P}_A(S_1)$  be the matrix for the orthogonal projection onto the subspace spanned by the eigenvectors of  $\mathbf{A}$  whose corresponding eigenvalues lies in  $S_1$ . Let  $\mathcal{P}_K$  be defined similarly. Then by the sin  $\Theta$  theorem [13] we have

$$\|\mathcal{P}_A(S_1) - \mathcal{P}_K(S_1)\| \leq \frac{\|\mathbf{A} - \mathbf{K}\|}{\text{dist}(S_1, S_2)}.$$

By equation (3.5) and Proposition 3.2, we have, with probability at least  $(1 - 2\eta)$ ,

$$\|\mathcal{P}_A(S_1) - \mathcal{P}_K(S_1)\| \leq \frac{2\sqrt{n \log(n/\eta)}}{n\delta_d - 4(1 + \sqrt{2})\sqrt{n \log(n/\eta)}} \leq 4\sqrt{\frac{\log(n/\eta)}{n\delta_d^2}},$$

provided that  $4(1 + \sqrt{2})\sqrt{n \log(n/\eta)} \leq n\delta_d/2$ .

To complete the proof, we note that if  $4(1 + \sqrt{2})\sqrt{n \log(n/\eta)} \leq n\delta_d/2$ , then  $S_1$  and  $S_2$  are disjoint. Thus  $\mathcal{P}_K(S_1) = \mathbf{U}_K \mathbf{U}_K^T$ . Finally, if  $\|\mathbf{A} - \mathbf{K}\| \leq 2\sqrt{n \log(n/\eta)}$ , then the eigenvalues of  $\mathbf{A}$  that lie in  $S_1$  are exactly the  $d$  largest eigenvalues of  $\mathbf{A}$  and  $\mathcal{P}_A(S_1) = \mathbf{U}_A \mathbf{U}_A^T$ . Equation (3.4) is thus established.  $\square$

Let  $\mathcal{H}$  be the reproducing kernel Hilbert space for  $\kappa$ . We now introduce a linear operator  $\mathcal{K}_{\mathcal{H},n}$  on  $\mathcal{H}$  defined as follows:

$$\mathcal{K}_{\mathcal{H},n}\eta = \frac{1}{n} \sum_{i=1}^n \langle \eta, \kappa(\cdot, X_i) \rangle_{\mathcal{H}} \kappa(\cdot, X_i).$$

The operator  $\mathcal{K}_{\mathcal{H},n}$  is the extension of  $\mathbf{K}$  as an operator on  $\mathbb{R}^n$  to an operator on  $\mathcal{H}$ . That is,  $\mathcal{K}_{\mathcal{H},n}$  is a linear operator on  $\mathcal{H}$  induced by  $\kappa$  and the  $X_1, X_2, \dots, X_n$ . The eigenvalues of  $\mathcal{K}_{\mathcal{H},n}$  and  $\mathbf{K}$  coincide, and furthermore, an

eigenfunction of  $\mathcal{K}_{\mathcal{H},n}$  is a linear interpolation of the corresponding eigenvector of  $\mathbf{K}$ . The reader is referred to Appendix B for more details.

The next result states that the rows of  $\mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}$  correspond to projecting the  $\Phi(X_i)$  using  $\hat{\mathcal{P}}_d$ , where  $\hat{\mathcal{P}}_d$  is the projection onto the  $d$ -dimensional subspace spanned by the eigenfunctions associated with the  $d$  largest eigenvalues of  $\mathcal{K}_{\mathcal{H},n}$ . We note that for large  $n$ ,  $\hat{\mathcal{P}}_d$  is close to the projection onto the  $d$ -dimensional subspace spanned by the eigenfunctions associated with the  $d$  largest eigenvalues of  $\mathcal{K}$  with high probability; see Theorem B.2.

LEMMA 3.4. *Let  $\hat{\mathcal{P}}_d$  be the projection onto the subspace spanned by the eigenfunctions corresponding to the  $d$  largest eigenvalues of  $\mathcal{K}_{\mathcal{H},n}$ . The rows of  $\mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}$  then correspond, up to some orthogonal transformation, to projections of the feature map  $\Phi$  onto  $\mathbb{R}^d$  via  $\hat{\mathcal{P}}_d$ , that is, there exists a unitary matrix  $\mathbf{W} \in \mathcal{M}_d(\mathbb{R})$  such that*

$$(3.6) \quad \mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}\mathbf{W} = [\iota(\hat{\mathcal{P}}_d(\Phi(X_1)))^T | \dots | \iota(\hat{\mathcal{P}}_d(\Phi(X_n)))^T]^T,$$

where  $\iota$  is the isometric isomorphism of a finite-dimensional Hilbert space onto  $\mathbb{R}^d$ .

The proof of Lemma 3.4 is given in the Appendix.

PROOF OF THEOREM 3.1. We first note that the sum of any row of  $\mathbf{A}$  is bounded from above by  $n$ , thus  $\|\mathbf{A}\| \leq n$ . Similarly,  $\|\mathbf{K}\| \leq n$ . On combining equation (3.4) and equation (3.3), we have, with probability at least  $1 - 2\eta$ ,

$$\begin{aligned} \|\mathcal{P}_{\mathbf{A}}\mathbf{A} - \mathcal{P}_{\mathbf{K}}\mathbf{K}\| &\leq \|\mathcal{P}_{\mathbf{A}}(\mathbf{A} - \mathbf{K})\| + \|(\mathcal{P}_{\mathbf{A}} - \mathcal{P}_{\mathbf{K}})\mathbf{K}\| \\ &\leq 2\sqrt{n \log(n/\eta)} + 4\delta_d^{-1}\sqrt{n \log(n/\eta)} \\ &\leq 6\delta_d^{-1}\sqrt{n \log(n/\eta)}. \end{aligned}$$

By Lemma A.1 in the Appendix, there exists an orthogonal  $\mathbf{W} \in \mathcal{M}_d(\mathbb{R})$  such that

$$\begin{aligned} \|\mathbf{U}_{\mathbf{A}}\mathbf{S}_{\mathbf{A}}^{1/2}\mathbf{W} - \mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}\| &\leq 6\delta_d^{-1}\sqrt{n \log(n/\eta)} \frac{\sqrt{d\|\mathcal{P}_{\mathbf{A}}\mathbf{A}\|} + \sqrt{d\|\mathcal{P}_{\mathbf{K}}\mathbf{K}\|}}{\lambda_d(\mathbf{K})} \\ &\leq 12\delta_d^{-1} \frac{n\sqrt{\log(n/\eta)}}{\lambda_d(\mathbf{K})}. \end{aligned}$$

We note that  $\lambda_d(\mathbf{K}) \geq n\lambda_d(\mathcal{K})/2$  provided that  $n$  satisfies  $\lambda_d(\mathcal{K}) > 4\sqrt{2} \times \sqrt{n^{-1} \log(n/\eta)}$ . Thus, we have

$$(3.7) \quad \|\mathbf{U}_{\mathbf{A}}\mathbf{S}_{\mathbf{A}}^{1/2}\mathbf{W} - \mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}\|_F \leq 24\delta_d^{-1} \frac{\sqrt{d \log(n/\eta)}}{\lambda_d(\mathcal{K})} \leq 24\delta_d^{-2} \sqrt{d \log(n/\eta)}$$

with probability at least  $1 - 2\eta$ .

Now, by Lemma 3.4, the rows of  $\mathbf{U}_K \mathbf{S}_K^{1/2}$  are (up to some orthogonal transformation) the projections of the feature map  $\Phi$  onto  $\mathbb{R}^d$  via  $\hat{\mathcal{P}}_d$ . On the other hand,  $\Phi_d(X)$  is the projection of  $\kappa(\cdot, X)$  onto  $\mathbb{R}^d$  via  $\mathcal{P}_d$ . By Theorem B.2 in the Appendix, for all  $X$ , we have

$$\|\hat{\mathcal{P}}_d \kappa(\cdot, X) - \mathcal{P}_d \kappa(\cdot, X)\|_{\mathcal{H}} \leq \|\hat{\mathcal{P}}_d - \mathcal{P}_d\|_{\text{HS}} \|\kappa(\cdot, X)\|_{\mathcal{H}} \leq 2\sqrt{2} \frac{\sqrt{\log(1/\eta)}}{\delta_d \sqrt{n}}$$

with probability at least  $1 - 2\eta$ . We therefore have, for some orthogonal  $\tilde{\mathbf{W}} \in \mathcal{M}_d(\mathbb{R})$ ,

$$(3.8) \quad \|\mathbf{U}_K \mathbf{S}_K^{1/2} \tilde{\mathbf{W}} - \Phi_d\|_F \leq 2\sqrt{2} \frac{\sqrt{\log(1/\eta)}}{\delta_d}$$

with probability at least  $1 - 2\eta$ . Equation (3.1) in the statement of the theorem then follows from equation (3.7) and equation (3.8).

To show equation (3.2), we first note that as the  $\{X_i\}_{i=1}^n$  are independent and identically distributed, the  $\{\hat{\Phi}_d(X_i)\}_{i=1}^n$  are exchangeable and hence identically distributed. Let  $\eta = n^{-2}$ . By conditioning on the event in equation (3.1), we have

$$(3.9) \quad \begin{aligned} \mathbb{E}[\|\hat{\Phi}_d(X_i) - \Phi_d(X_i)\|] &\leq \sqrt{\mathbb{E}[\|\hat{\Phi}_d(X_i) - \Phi_d(X_i)\|^2]} \\ &\leq \sqrt{\frac{1}{n} \mathbb{E}[\|\hat{\Phi}_d - \Phi_d\|_F^2]} \\ &\leq \frac{1}{\sqrt{n}} \sqrt{\left(1 - \frac{2}{n^2}\right) (27\delta_d^{-2} \sqrt{3d \log n})^2 + \frac{2}{n^2} 2n} \\ &\leq 27\delta_d^{-2} \sqrt{\frac{6d \log n}{n}}, \end{aligned}$$

because the worst case bound is  $\|\hat{\Phi}_d - \Phi_d\|_F \leq 2n$  with probability 1. Equation (3.2) follows from equation (3.9) and Markov’s inequality.  $\square$

**4. Universally consistent vertex classification.** The results in Section 3 show that by using the eigen-decomposition of  $\mathbf{A}$ , we can consistently estimate the truncated feature map  $\Phi_d$  for any fixed, finite  $d$  (up to an orthogonal transformation). In the subsequent discussion, we will often refer to the rows of the eigen-decomposition of  $\mathbf{A}$ , that is, the rows of  $\mathbf{U}_A \mathbf{S}_A^{1/2}$  as the *estimated vectors*. Sussman, Tang and Priebe [36] showed that, for the dot product kernel on a finite-dimensional space  $\mathcal{X}$ , the  $k$ -nearest-neighbors classifier on  $\mathbb{R}^d$  is universally consistent when we select the neighbors using the estimated vectors rather than the true but unknown latent positions. This result can be trivially extended to the setting for an arbitrary finite-rank kernel  $\kappa$  as long as the feature map  $\Phi$  of  $\kappa$  is

injective. It is also easy to see that if the feature map  $\Phi$  is not injective, then any classifier that uses only the estimated vectors (or the feature map  $\Phi$ ) is no longer universally consistent. This section is concerned with the setting where the kernel  $\kappa$  is an infinite-rank kernel with an injective feature map  $\Phi$  onto  $l_2$ . Well-known examples of these kernels are the class of universal kernels [25, 30, 31].

DEFINITION 4.1. A continuous kernel  $\kappa$  on some metric space  $(\mathcal{X}, d)$  is a universal kernel if for some feature map  $\Phi : \mathcal{X} \mapsto H$  of  $\kappa$  to some Hilbert space  $H$ , the class of functions of the form

$$\mathcal{F}_\Phi = \{ \langle w, \Phi \rangle_H : w \in H \}$$

is dense in  $\mathcal{C}(\mathcal{X})$ ; that is, for any continuous function  $g : \mathcal{X} \mapsto \mathbb{R}$  and any  $\varepsilon > 0$ , there exists a  $f \in \mathcal{F}_\Phi$  such that  $\|f - g\|_\infty < \varepsilon$ .

We note that if  $\mathcal{F}_\Phi$  is dense in  $\mathcal{C}(\mathcal{X})$  for some feature map  $\Phi$  and  $\Phi' : \mathcal{X} \mapsto H'$  is another feature map of  $\kappa$ , then  $\mathcal{F}_{\Phi'}$  is also dense in  $\mathcal{C}(\mathcal{X})$ , that is, the universality of  $\kappa$  is independent of the choice for its feature map. Furthermore, every feature map of a universal kernel is injective.

The following result lists several well-known universal kernels.

PROPOSITION 4.2 ([25, 31]). *Let  $S$  be a compact subset of  $\mathbb{R}^d$ . Then the following kernels are universal on  $S$ :*

- the exponential kernel  $\kappa(x, y) = \exp(\langle x, y \rangle)$ ;
- the Gaussian kernel  $\kappa(x, y) = \exp(-\|x - y\|^2 / \sigma^2)$  for all  $\sigma > 0$ ;
- the binomial kernel  $\kappa(x, y) = (1 - \langle x, y \rangle)^{-\alpha}$  for  $\alpha > 0$ ;
- the inverse multiquadrics  $\kappa(x, y) = (c^2 + \|x - y\|^2)^{-\beta}$  with  $c > 0$  and  $\beta > 0$ .

If the kernel matrix  $\mathbf{K}$  is known, then results on the universal consistency of support vector machines with universal kernels are available; see, for example, [32]. If the feature map  $\Phi$  is known, then Biau, Bunea and Wegkamp [5] showed that the  $k$ -nearest-neighbors on  $\Phi_d$  are universally consistent as  $k \rightarrow \infty$  and  $d \rightarrow \infty$  where  $k$  and  $d$  are chosen using a structural risk minimization approach.

Our universally consistent classifier operates on the estimated vectors and is based on an empirical risk minimization approach. Namely, we will show that the classifier that minimizes a convex surrogate  $\varphi$  for 0–1 loss from a class of linear classifiers  $\mathcal{C}^{(d_n)}$  is universally consistent provided that the convex surrogate  $\varphi$  satisfies some mild conditions and that the complexity of the class  $\mathcal{C}^{(d_n)}$  grows in a controlled manner.

First, we will expand our framework to the classification setting. Let  $\mathcal{X}$  be as in Section 2, and let  $F_{\mathcal{X}, Y}$  be a distribution on  $\mathcal{X} \times \{-1, 1\}$ . Let  $(X_1, Y_1), \dots, (X_{n+1}, Y_{n+1}) \stackrel{\text{i.i.d.}}{\sim} F_{\mathcal{X}, Y}$ , and let  $\mathbf{K}$  and  $\mathbf{A}$  be as in Section 2. The  $\{Y_i\}$  are the class labels for the vertices in the graph corresponding to the adjacency matrix  $\mathbf{A}$ .

We suppose that we observe only  $\mathbf{A}$ , the adjacency matrix, and  $Y_1, \dots, Y_n$ , the class labels for all but the last vertex. Our goal is to accurately classify this last vertex, so for convenience of notation we shall define  $X := X_{n+1}$  and  $Y := Y_{n+1}$ . Let the rows of  $\mathbf{U}_\mathbf{A} \mathbf{S}_\mathbf{A}^{1/2}$  be denoted by  $\zeta_d(X_1), \dots, \zeta_d(X_{n+1})$  (even though the  $X_i$  are unobserved/unknown). We want to find a classifier  $h_n$  such that, for any distribution  $F_{\mathcal{X}, Y}$ ,

$$\begin{aligned} \mathbb{E}[L_n] &:= \mathbb{E}[\mathbb{P}[h_n(\zeta_{d_n}(X)) \neq Y | (\zeta_{d_n}(X_1), Y_1), \dots, (\zeta_{d_n}(X_n), Y_n)]] \\ &\rightarrow \mathbb{P}[h^*(X) \neq Y] =: L^*, \end{aligned}$$

where  $h^*$  is the Bayes-optimal classifier, and  $L^*$  is its associated Bayes-risk.

Let  $\mathcal{C}^{(d)}$  be the class of linear classifiers using the truncated feature map  $\Phi_d$  whose linear coefficients are normalized to have norm at most  $d$ , that is,  $g \in \mathcal{C}^{(d)}$ , if and only if  $g$  is of the form

$$(4.1) \quad g(x) = \begin{cases} 1, & \text{if } \langle w, \Phi_d(x) \rangle > 0, \\ -1, & \text{if } \langle w, \Phi_d(x) \rangle \leq 0, \end{cases}$$

for some  $w \in \mathbb{R}^d$  with  $\|w\| \leq d$ . We note that the  $\{\mathcal{C}^{(d)}\}$  are increasing, that is,  $\mathcal{C}^{(d)} \subset \mathcal{C}^{(d')}$  for  $d < d'$  and that  $\bigcup_{d \geq 1} \mathcal{C}^{(d)} = \mathcal{F}_\Phi = \{\langle w, \Phi \rangle_{\mathcal{H}} : w \in \mathcal{H}\}$ . Because  $\kappa$  is universal,  $\mathcal{F}_\Phi$  is dense in  $\mathcal{C}(\mathcal{X})$  and as  $\mathcal{X}$  is compact,  $\mathcal{F}_\Phi$  is dense in the space of measurable functions on  $\mathcal{X}$ . Thus  $\lim_{d \rightarrow \infty} \inf_{g \in \mathcal{C}^{(d)}} L(g) = L^*$  and so one can show that empirical risk minimization over the class  $\mathcal{C}^{(d_n)}$  for any increasing and divergent sequence  $(d_n)$  yields a universally consistent classifier (Theorem 18.1 in [14]). The remaining part of this section is concerned with modifying this result so that it applies to the estimated feature map  $\zeta_d$  instead of the true feature map  $\Phi_d$ .

We now describe a setup for empirical risk minimization over  $\mathcal{C}^{(d)}$  for increasing  $d$  where we use the estimated  $\zeta_d$  in place of the  $\Phi_d$ . Let us write  $\hat{L}_n(w; \zeta_d)$  for the empirical error when using the  $\zeta_d$ , that is,

$$\hat{L}_n(w; \zeta_d) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{\text{sign}(\langle w, \zeta_d(X_i) \rangle) \neq Y_i\} \leq \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Y_i \langle w, \zeta_d(X_i) \rangle \leq 0\}.$$

We want to show that minimization of  $\hat{L}_n(w; \zeta_{d_n})$  over the class  $\mathcal{C}^{(d_n)}$  for increasing  $(d_n)$  leads to a universally consistent classifier for our latent position graphs setting. However, the loss function  $L(f) = \mathbb{P}(\text{sign}(f(X)) \neq Y)$  of a classifier  $f$  as well as its empirical version  $\hat{L}_n(f)$  is based on the 0–1 loss, which is discontinuous at  $f = 0$ . Furthermore, the distribution of  $\zeta_d(X)$  not available. This induces complications in relating  $\hat{L}_n(w; \zeta_d)$  to  $\hat{L}_n(w; \Phi_d)$ . That is, the classifier obtained by minimizing the 0–1 loss using  $\zeta$  might be very different from the classifier obtained by minimizing the 0–1 loss using  $\Phi$ .

To circumvent this issue, we will work with some convex loss function  $\varphi$  that is a surrogate of the 0–1 loss. The notion of constructing classification algorithms that correspond to minimization of a convex surrogate for the 0–1 loss is a powerful

one and the authors of [2, 24, 40], among others, showed that one can obtain, under appropriate regularity conditions, Bayes-risk consistent classifiers in this manner.

Let  $\varphi : \mathbb{R} \mapsto [0, \infty)$ . We define the  $\varphi$ -risk of  $f : \mathcal{X} \mapsto \mathbb{R}$  by

$$R_\varphi(f) = \mathbb{E}\varphi(Yf(X)).$$

Given some data  $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$ , the empirical  $\varphi$ -risk of  $f$  is defined as

$$\hat{R}_{\varphi,n}(f) = \frac{1}{n} \sum_{i=1}^n \varphi(Y_i f(X_i)).$$

We will often write  $\hat{R}_\varphi(f)$  if the number of samples  $(X_i, Y_i)$  in  $\mathcal{D}_n$  is clear from the context. Let  $w \in \mathbb{R}^d, \|w\| \leq d$  index a linear classifier on  $\mathcal{C}^{(d)}$ . Denote by  $R_\varphi(w; \Phi_d), \hat{R}_{\varphi,n}(w; \Phi_d), \hat{R}_{\varphi,n}(w; \zeta_d)$  and  $R_{\varphi,n}(w; \zeta_d)$  the various quantities analogous to  $L(w; \Phi_d), \hat{L}_n(w; \Phi_d), \hat{L}_n(w; \zeta_d)$  and  $L_n(w; \zeta_d)$  for 0–1 loss defined previously. Let us also define  $R_\varphi^*$  as the minimum  $\varphi$ -risk over all measurable functions  $f : \mathcal{X} \mapsto \mathbb{R}$ .

In this paper, we will assume that the convex surrogate  $\varphi : \mathbb{R} \mapsto [0, \infty)$  is differentiable with  $\varphi'(0) < 0$ . This implies that  $\varphi$  is *classification-calibrated* [2]. Examples of classification-calibrated loss functions are the exponential loss function  $\varphi(x) = \exp(-x)$  in boosting, the logit function  $\varphi(x) = \log_2(1 + \exp(-x))$  in logistic regression and the square error loss  $\varphi(x) = (1 - x)^2$ . For classification-calibrated loss functions, we have the following result.

**THEOREM 4.3 ([2]).** *Let  $\varphi : \mathbb{R} \mapsto [0, \infty)$  be classification-calibrated. Then for any sequence of measurable functions  $f_i : \mathcal{X} \mapsto \mathbb{R}$  and every probability distribution  $F_{\mathcal{X},\mathcal{Y}}, R_\varphi(f_i) \rightarrow R_\varphi^*$  implies  $L(f_i) \rightarrow L^*$ .*

We now state the main result of this section, which is that empirical  $\varphi$ -risk minimization over the class  $\mathcal{C}^{(d_n)}$  for some diverging sequence  $(d_n)$  yields a universally consistent classifier for the latent position graphs setting.

**THEOREM 4.4.** *Let  $\varepsilon \in (0, 1/4)$  be fixed. For a given  $d$ , let  $C_d = \max\{\varphi'(-d), \varphi'(d)\}$ . Suppose that  $d_n$  is given by the following rule:*

$$(4.2) \quad d_n = \max \left\{ d \leq n : \frac{1}{n}(\lambda_d(\mathbf{A}) - \lambda_{d+1}(\mathbf{A})) \geq 32\sqrt{d}C_d \left( \frac{d \log n}{n} \right)^{1/4-\varepsilon} \right\}.$$

*Let  $\tilde{g}_n$  be the classifier obtained by empirical  $\varphi$ -risk minimization over  $\mathcal{C}^{(d_n)}$ . Then  $R_{\varphi,n}(\tilde{g}_n) \rightarrow R_\varphi^*$  as  $n \rightarrow \infty$  and  $\text{sign}(\tilde{g}_n)$  is universally consistent, that is,*

$$\mathbb{E}[\mathbb{P}(\text{sign}(\tilde{g}_n(\zeta_{d_n}(X))) \neq Y | \mathcal{D}_n)] \rightarrow L^*$$

*as  $n \rightarrow \infty$  for any distribution  $F_{\mathcal{X},\mathcal{Y}}$ .*

REMARK. We note that due to the use of the estimated  $\zeta$  in place of the true  $\Phi$ , Theorem 4.4 is limited in two key aspects. The first is that we do not claim that  $\tilde{g}_n$  is universally *strongly* consistent for any  $F_{\mathcal{X},Y}$  and the second is that we cannot specify  $d_n$  in advance. In return, the minimization of the empirical  $\varphi$ -risk over the class  $\mathcal{C}^{(d)}$  is a convex optimization problem and the solution can be obtained more readily than the minimization of empirical 0–1 loss. For example, by using squared error loss instead of 0–1 loss, the classifier that minimizes the empirical  $\varphi$ -risk can be viewed as a ridge regression problem. We note also that as the only accumulation point in the spectrum of  $\mathcal{K}$  is at zero, the sequence  $(d_n)$  as specified in equation (4.2) exists. Furthermore, such a sequence is only one possibility among many. In particular, the conclusion in Theorem 4.4 holds for any sequence  $(d_n)$  that diverges and satisfies the condition  $\delta_{d_n}^2 = o(n^{-1/2}d^{3/2}\sqrt{\log n})$ . Choosing the right  $(d_n)$  requires balancing the approximation error  $\inf_{g \in \mathcal{C}^{(d_n)}} R_\varphi(g) - R_\varphi^*$  and the estimation error  $R_\varphi(\tilde{g}_n) - \inf_{w \in \mathcal{C}^{(d_n)}} R_\varphi(g)$ , and this can be done using an approach based on structural risk minimization; see, for example, Section 18.1 of [14] and [24].

We now proceed to prove Theorem 4.4. A rough sketch of the argument goes as follows. First we show that any classifier  $g$  using the estimated vectors  $\zeta_d$  induces a classifier  $g'$  using the true truncated feature map  $\Phi_d$  such that the empirical  $\varphi$ -risk of  $g$  is “close” to the empirical  $\varphi$ -risk of  $g'$ . Then by applying a Vapnik–Chervonenkis-type bound for  $g'$ , we show that the classifier  $\tilde{g}$  (using  $\zeta_d$ ) selected by empirical  $\varphi$ -risk minimization induces a classifier  $\hat{g}$  (using  $\Phi_d$ ) with the  $\varphi$ -risk of  $\hat{g}$  being “close” to the minimum  $\varphi$ -risk for the classifiers in the class  $\mathcal{C}^{(d)}$ . Universal consistency of  $\hat{g}$  and hence of  $\tilde{g}$  follows by letting  $d$  grow in a specified manner.

Let  $1 \leq d \leq n$ . Let  $\mathbf{U}_A \mathbf{S}_A^{1/2}$  be the embedding of  $\mathbf{A}$  into  $\mathbb{R}^d$ . Let  $\mathbf{W}_d \in \mathcal{M}_d(\mathbb{R})$  be an orthogonal matrix given by

$$\mathbf{W}_d = \min_{\mathbf{W}: \mathbf{W}^T \mathbf{W} = \mathbf{I}} \|\mathbf{U}_A \mathbf{S}_A^{1/2} \mathbf{W} - \Phi_d\|_F.$$

The following result states that if there is a gap in the spectrum of  $\mathcal{K}$  at  $\lambda_d(\mathcal{K})$ , then  $\hat{R}_{\varphi,n}(w; \zeta_d)$  and  $\hat{R}_{\varphi,n}(\mathbf{W}_d w, \Phi_d)$  is close for all  $w \in \mathbb{R}^d$ ,  $\|w\| \leq d$ . That is, the empirical  $\varphi$ -risk of a linear classifier using  $\zeta_d$  is not too different from the empirical  $\varphi$ -risk of a related classifier (the relationship is given by  $\mathbf{W}_d$ ) using  $\Phi_d$ .

PROPOSITION 4.5. *Let  $d \geq 1$  be such that  $\lambda_d(\mathcal{K}) > \lambda_{d+1}(\mathcal{K})$ , and let  $C_d = \max\{\varphi'(d), \varphi'(-d)\}$ . Then for any  $w \in \mathbb{R}^d$ ,  $\|w\| \leq d$ , we have, with probability at least  $1 - 1/n^2$ ,*

$$(4.3) \quad \left| \hat{R}_{\varphi,n}(w; \zeta_d) - \hat{R}_{\varphi,n}(\mathbf{W}_d w; \Phi_d) \right| \leq 27\delta_d^{-2} d C_d \sqrt{\frac{3d \log n}{n}}.$$

PROOF. We have

$$\begin{aligned} & \hat{R}_{\varphi,n}(w; \zeta_d) - \hat{R}_{\varphi,n}(\mathbf{W}^d w; \Phi_d) \\ &= \frac{1}{n} \sum_{i=1}^n \varphi(Y_i \langle w, \zeta_d(X_i) \rangle) - \varphi(Y_i \langle \mathbf{W}^d w, \Phi_d(X_i) \rangle). \end{aligned}$$

Now  $\varphi$  is convex and thus locally Lipschitz-continuous. Also,  $|\langle w, \Phi_d(X) \rangle| \leq d$  for all  $X \in \mathcal{X}$ . Hence, there exists a constant  $M$  independent of  $n$  and  $F_{\mathcal{X},\mathcal{Y}}$  such that

$$\begin{aligned} & |\varphi(Y_i \langle w, \zeta_d(X_i) \rangle) - \varphi(Y_i \langle \mathbf{W}_d w, \Phi_d(X_i) \rangle)| \\ & \leq M \left| Y_i \left\langle \frac{w}{\|w\|}, \zeta_d(X_i) \right\rangle - Y_i \left\langle \mathbf{W}_d \frac{w}{\|w\|}, \Phi_d(X_i) \right\rangle \right| \end{aligned}$$

for all  $i$ . Thus, by Theorem 3.1, we have

$$\begin{aligned} & |\hat{R}_{\varphi,n}(w; \zeta_d) - \hat{R}_{\varphi,n}(\mathbf{W}_d w; \Phi_d)| \\ & \leq \frac{M}{n} \sum_{i=1}^n \left| Y_i \left\langle \frac{w}{\|w\|}, \zeta_d(X_i) \right\rangle - Y_i \left\langle \mathbf{W}_d \frac{w}{\|w\|}, \Phi_d(X_i) \right\rangle \right| \\ & \leq \frac{M}{n} \sum_{i=1}^n \|\zeta_d(X_i) - (\mathbf{W}_d)^T \Phi_d(X_i)\| \\ (4.4) \quad & \leq \frac{M}{\sqrt{n}} \left( \sum_{i=1}^n \|\zeta_d(X_i) - (\mathbf{W}_d)^T \Phi_d(X_i)\|^2 \right)^{1/2} \\ & \leq \frac{M}{\sqrt{n}} \|\mathbf{U}_A \mathbf{S}_A^{1/2} - (\mathbf{W}_d)^T \Phi_d\|_F \\ & \leq 27\delta_d^{-2} M \sqrt{\frac{3d \log n}{n}} \end{aligned}$$

with probability at least  $1 - 1/n^2$ . By the mean-value theorem, we can take  $M = d \max\{\varphi'(d), \varphi'(-d)\}$  to complete the proof.  $\square$

The Vapnik–Chervonenkis theory for 0–1 loss function can also be extended to the convex surrogate setting  $R_\varphi$ . In particular, the following result provides a uniform deviation bound for  $|R_\varphi(f) - \hat{R}_{\varphi,n}(f)|$  for functions  $f$  in some class  $\mathcal{F}$  in terms of the VC-dimension of  $\mathcal{F}$ .

LEMMA 4.6 ([24]). *Let  $\mathcal{F}$  be a class of functions with VC-dimension  $V < \infty$ . Suppose that the range of any  $f \in \mathcal{F}$  is contained in the interval  $[-d, d]$ . Let  $n \geq 5$ .*

Then we have, with probability at least  $1 - 1/n^2$ ,

$$(4.5) \quad \sup_{f \in \mathcal{F}} |R_\varphi(f) - \hat{R}_{\varphi,n}(f)| \leq 10d \max\{\varphi'(d), \varphi'(-d)\} \sqrt{\frac{3V \log n}{n}}.$$

The following result combines Proposition 4.5 and Lemma 4.6 and shows that minimizing  $\hat{R}_{\varphi,n}(w; \zeta_d)$  over  $w \in \mathbb{R}^d, \|w\| \leq d$  leads to a classifier whose  $\varphi$ -risk is close to optimal in the class  $\mathcal{C}^{(d)}$  with high probability.

LEMMA 4.7. *Let  $d \geq 1$  be such that  $\lambda_d(\mathcal{K}) > \lambda_{d+1}(\mathcal{K})$  and let  $C_d = \max\{\varphi'(d), \varphi'(-d)\}$ . Let  $\tilde{w}_d$  minimize  $\hat{R}_{\varphi,n}(w; \zeta_d)$  over  $\mathbb{R}^d, \|w\| \leq d$ . Then with probability at least  $1 - 2/n^2$ ,*

$$(4.6) \quad R_\varphi(\mathbf{W}_d \tilde{w}_d; \Phi_d) - \inf_{w \in \mathcal{C}^{(d)}} R_\varphi(w; \Phi_d) \leq 74\delta_d^{-2} d C_d \sqrt{\frac{3d \log n}{n}}.$$

PROOF. For ease of notation, we let  $\varepsilon(n, d)$  be the term in the right-hand side of equation (4.3), and let  $C(n, d)$  be the term in the right-hand side of equation (4.5). Also let  $\bar{w}^{(d)} := \arg \inf_{w \in \mathcal{C}^{(d)}} R_\varphi(w; \Phi_d)$ . We then have

$$\begin{aligned} R_\varphi(\mathbf{W}_d \tilde{w}_d; \Phi_d) &\leq \hat{R}_{\varphi,n}(\mathbf{W}_d \tilde{w}_d; \Phi_d) + C(n, d) \\ &\leq \hat{R}_{\varphi,n}(\tilde{w}_d; \zeta_d) + \varepsilon(n, d) + C(n, d) \\ &\leq \hat{R}_{\varphi,n}((\mathbf{W}_d)^T \tilde{w}_d; \zeta_d) + \varepsilon(n, d) + C(n, d) \\ &\leq \hat{R}_{\varphi,n}(\bar{w}_d; \Phi_d) + 2\varepsilon(n, d) + C(n, d) \\ &\leq R_\varphi(\bar{w}_d; \Phi_d) + 2\varepsilon(n, d) + 2C(n, d) \end{aligned}$$

with probability at least  $1 - 2/n^2$ .  $\square$

REMARK. Equation (4.6) is a VC-type bound. The term  $d^{3/2}\delta_d^{-2}$  in equation (4.6) can be viewed as contributing to the generalization error for the classifiers in  $\mathcal{C}^{(d)}$ . That is, because we are training using the estimated vectors in  $\mathbb{R}^d$ , the generalization error not only depends on the dimension of the embedded space, but also depends on how accurate the estimated vectors are in that space.

We now have the necessary ingredients to prove the main result of this section.

PROOF OF THEOREM 4.4. Let  $(d_n)$  be a nondecreasing sequence of positive integers that diverges to  $\infty$  and that

$$(4.7) \quad \delta_{d_n}^{-2} d_n C_{d_n} \sqrt{\frac{d \log n}{n}} = o(1).$$

By Lemma 4.7 and the Borel–Cantelli lemma, we have

$$\lim_{n \rightarrow \infty} \left[ R_\varphi(\mathbf{W}_{d_n} \tilde{w}_{d_n}; \Phi_{d_n}) - \inf_{w \in \mathcal{C}^{(d_n)}} R_\varphi(w; \Phi_{d_n}) \right] = 0$$

almost surely. As  $(d_n)$  diverges,  $\lim_{n \rightarrow \infty} \inf_{w \in \mathcal{C}^{(d_n)}} R_\varphi(w; \Phi_{d_n}) = R_\varphi^*$  by Proposition A.2. We therefore have

$$\lim_{n \rightarrow \infty} R_\varphi(\mathbf{W}_{d_n} \tilde{w}_{d_n}; \Phi_{d_n}) = R_\varphi^*$$

almost surely. Now fix a  $n$ . The empirical  $\varphi$ -risk minimization on  $w \in \mathbb{R}^{d_n}$ ,  $\|w\| \leq d_n$  using the estimated vectors  $\zeta_{d_n}$  gives us a classifier  $\langle \tilde{w}_{d_n}, \zeta_{d_n} \rangle$ . We now consider the difference  $R_{\varphi,n}(\tilde{w}_{d_n}; \zeta_{d_n}) - R_\varphi(\mathbf{W}_{d_n} \tilde{w}_{d_n}; \Phi_{d_n})$ . By a similar computation to that used in the derivation of equation (4.4), we have

$$\begin{aligned} & R_{\varphi,n}(\tilde{w}_{d_n}; \zeta_{d_n}) - R_\varphi(\mathbf{W}_{d_n} \tilde{w}_{d_n}; \Phi_{d_n}) \\ &= \left| \mathbb{E}[\varphi(Y \langle \tilde{w}_{d_n}, \zeta_{d_n}(X) \rangle)] - \mathbb{E}[\varphi(Y \langle \mathbf{W}_{d_n} \tilde{w}_{d_n}, \Phi_{d_n}(X) \rangle)] \right| \\ &\leq d_n C_{d_n} \mathbb{E}[\|\zeta_{d_n}(X) - (\mathbf{W}_{d_n})^T \Phi_{d_n}(X)\|] \\ &\leq d_n C_{d_n} \sqrt{\mathbb{E}[\|\zeta_{d_n}(X) - (\mathbf{W}_{d_n})^T \Phi_{d_n}(X)\|^2]} \\ &\leq 27\delta_{d_n}^{-2} d_n C_{d_n} \sqrt{\frac{6d \log n}{n}} = o(1). \end{aligned}$$

We therefore have

$$\lim_{n \rightarrow \infty} R_{\varphi,n}(\tilde{w}_{d_n}; \zeta_{d_n}) = R_\varphi^*.$$

Thus, by Theorem 4.3, we have

$$\lim_{n \rightarrow \infty} \mathbb{E}[L_n(\tilde{w}_{d_n}; \zeta_{d_n})] = L^*.$$

The only thing that remains is the use of  $\frac{1}{n}(\lambda_d(\mathbf{A}) - \lambda_{d+1}(\mathbf{A}))$  as an estimate for  $\delta_d$ . By Proposition 3.2 and Theorem 3.1, we have

$$(4.8) \quad \sup_{d \geq 1} \left| \delta_d - \frac{1}{n}(\lambda_d(\mathbf{A}) - \lambda_{d+1}(\mathbf{A})) \right| \leq 10 \sqrt{\frac{\log(2/n^2)}{n}}$$

with probability at least  $1 - 2/n^2$ . Thus, if  $d_n$  satisfy equation (4.2), then equation (4.8) implies that equation (4.7) holds for  $d_n \rightarrow \infty$  with probability at least  $1 - 2/n^2$ . Finally, we note that as  $n \rightarrow \infty$ , there exists a sequence  $(d_n)$  that satisfies equation (4.2) and diverges to  $\infty$ , as the only accumulation point in the spectrum of  $\mathcal{H}$  is at zero.  $\square$

**5. Conclusions.** In this paper we investigated the problem of finding a universally consistent classifier for classifying the vertices of latent position graphs. We showed that if the link function  $\kappa$  used in the construction of the graphs belong to the class of universal kernels, then an empirical  $\varphi$ -risk minimization approach, that is, minimizing a convex surrogate of the 0–1 loss over the class of linear classifiers in  $\mathbb{R}^{d_n}$  for some sequence  $d_n \rightarrow \infty$ , yields universally consistent vertices classifiers.

We have presented the universally consistent classifiers in the setting where the graphs are on  $n + 1$  vertices, there are  $n$  labeled vertices and the task is to classify the remaining unlabeled vertex. It is easy to see that in the case where there are only  $m < n$  labeled vertices, the same procedure given in Theorem 4.4 with  $n$  replaced by  $m$  still yields universally consistent classifiers, provided that  $m \rightarrow \infty$ .

The bound for the generalization error of the classifiers in Section 4 is of the form  $O(n^{-1/2} \delta_d^{-2} \sqrt{d^3 \log n})$ . This bound depends on both the subspace projection error in Section 3 as well as the generalization error of the class  $\mathcal{C}^{(d_n)}$ . It is often the case that the bound on the generalization error of the class  $\mathcal{C}^{(d_n)}$  can be improved, as long as the classification problems satisfy a “low-noise” condition, that is, that the posterior probability  $\eta(x) = \mathbb{P}[Y = 1|X = x]$  is bounded away from 1/2. Results on fast convergence rates in low-noise conditions, for example, [2, 6] can thus be used, but as the subspace projection error is independent of the low-noise condition, there might not be much improvement in the resulting error bound.

Also related to the above issue is the choice of the sequence  $(d_n)$ . If more is known about the kernel  $\kappa$ , then the choice for the sequence  $(d_n)$  can be adjusted accordingly. For example, good bounds for  $\Lambda_k = \sum_{j \geq k} \lambda_j(\mathcal{K})$ , the sum of the tail eigenvalues of  $\mathcal{K}$ , along with bounds for the error between the truncated feature map  $\Phi_d$  and the feature map  $\Phi$  from [7, 9, 29] can be used to select the sequence  $(d_n)$ .

The results presented in Section 3 and Section 4 implicitly assumed that the graphs arising from the latent position model are dense. It is possible to extend these results to sparse graphs. A sketch of the ideas is as follows. Let  $\rho_n \in (0, 1)$  be a scaling parameter, and consider the latent position model with kernel  $\kappa$  and distribution  $F$  for the latent features  $\{X_i\}$ . Given  $\{X_i\}_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} F$ , let  $\mathbf{K}_n = (\rho_n \kappa(X_i, X_j))_{i,j=1}^n$ , that is, the entries of  $\mathbf{K}_n$  are given by the kernel  $\kappa$  scaled by the scaling parameter  $\rho_n$ . This variant of the latent position model is also present in the notion of inhomogeneous random graphs [8, 26]. Given  $\mathbf{K}_n$ ,  $\mathbf{A}_n = \text{Bernoulli}(\mathbf{K}_n)$  is the adjacency matrix. The factor  $\rho_n$  controls the sparsity of the resulting latent position graph. For example,  $\rho_n = (\log n)/n$  leads to sparse, connected graphs almost surely while  $\rho_n = 1/n$  leads to graphs with a single giant connected component [8]. Suppose now that  $\rho_n = \Omega((\log n)/n)$ . The following result is a restatement of Theorem 3.1 for the latent position model in the presence of the scaling parameter  $\rho_n$ . Its proof is almost identical to that of Theorem 3.1

provided that one uses the bound in term of the maximum degree  $\Delta$  in Proposition 3.2. We note that  $\delta_d$  is defined in terms of the spectrum of  $\mathcal{K}$  which does not depend on the scaling parameter  $\rho_n$ , and similarly for the feature map  $\Phi$  and its truncation  $\Phi_d$ .

**THEOREM 5.1.** *Let  $d \geq 1$  be given. Denote by  $\delta_d$  the quantity  $\lambda_d(\mathcal{K}) - \lambda_{d+1}(\mathcal{K})$ , and suppose that  $\delta_d > 0$ . Then with probability greater than  $1 - 2\eta$ , there exists a unitary matrix  $\mathbf{W} \in \mathcal{M}_d(\mathbb{R})$  such that*

$$(5.1) \quad \|\rho_n^{-1/2} \mathbf{U}_A \mathbf{S}_A^{1/2} \mathbf{W} - \Phi_d\|_F \leq 27\delta_d^{-2} \sqrt{\frac{d \log(n/\eta)}{\rho_n}},$$

where  $\Phi_d$  denotes the matrix in  $\mathcal{M}_{n,d}(\mathbb{R})$  whose  $i$ th row is  $\Phi_d(X_i)$ . Let us denote by  $\hat{\Phi}_d(X_i)$  the  $i$ th row of  $\rho_n^{-1/2} \mathbf{U}_A \mathbf{S}_A^{1/2} \mathbf{W}$ . Then, for each  $i \in [n]$  and any  $\varepsilon > 0$ ,

$$(5.2) \quad \mathbb{P}[\|\hat{\Phi}_d(X_i) - \Phi_d(X_i)\| > \varepsilon] \leq 27\delta_d^{-2} \varepsilon^{-1} \sqrt{\frac{6d \log n}{n\rho_n}}.$$

Thus, for  $\rho_n = n^{-1+\varepsilon}$  for some  $\varepsilon > 0$  [or even  $\rho_n = (\log^k n)/n$  for some sufficient large  $k$ ], equation (5.2) states that with high probability, the estimated feature map is (after scaling by  $\rho_n^{-1/2}$  and rotation) converging to the true truncated feature map  $\Phi_d$  as  $n \rightarrow \infty$ . The results from Section 4 can then be modified to show the existence of a universally consistent linear classifier. The main difference between the sparse setting and the dense setting would be the generalization bounds in Proposition 4.5 and Lemma 4.7. This would lead to a different selection rule for the sequence of embedding dimensions  $d_n \rightarrow \infty$  then the one in Theorem 5.1, that is, the  $d_n$  would diverge more slowly for the sparse setting compared to the dense setting. A precise statement and formulation of the results in Section 4 for the sparse setting might require some care, but should be for the most part straightforward. We also note that even though  $\rho_n$  is most likely unknown, one can scale the embedding  $\mathbf{Z} = \mathbf{U}_A \mathbf{S}_A^{1/2}$  by any value  $c_n$  that is of the same order as  $\rho_n^{-1/2}$ . An appropriate value for  $c_n$  is, for example, one that makes  $\max_i \|c_n Z_i\|^2 = 1$  where  $Z_i$  is the  $i$ th row of  $\mathbf{Z}$ .

Finally, we note it is of potential interest to extend the results herein to graphs with attributes on the edges, latent position graphs with nonpositive definite link functions  $\kappa$  and graphs with errorfully observed edges.

## APPENDIX A: ADDITIONAL PROOFS

**PROOF OF LEMMA 3.4.** Let  $\Psi_{r,n} \in \mathbb{R}^n$  be the vector whose entries are  $\sqrt{\lambda_r} \psi_r(X_i)$  for  $i = 1, 2, \dots, n$  with  $\lambda_r = \lambda_r(\mathcal{K})$ . We note that  $\mathbf{K} = \sum_{r=1}^{\infty} \Psi_{r,n} \times$

$\Psi_{r,n}^T$ . Let  $\hat{u}^{(1)}, \dots, \hat{u}^{(d)}$  be the eigenvectors associated with the  $d$  largest eigenvalues of  $\mathbf{K}/n$ . We have

$$\mathcal{P}_{\mathbf{K}}\mathbf{K} = \sum_{s=1}^d \sum_{r=1}^{\infty} \hat{u}^{(s)} (\hat{u}^{(s)})^T \Psi_{r,n} \Psi_{r,n}^T \hat{u}^{(s)} (\hat{u}^{(s)})^T.$$

The  $ij$ th entry of  $\mathcal{P}_{\mathbf{K}}\mathbf{K}$  is then given by

$$\sum_{s=1}^d \sum_{r=1}^{\infty} \hat{u}_i^{(s)} (\hat{u}^{(s)})^T \Psi_{r,n} \Psi_{r,n}^T \hat{u}^{(s)} \hat{u}_j^{(s)}.$$

Let  $\hat{v}^{(1)}, \dots, \hat{v}^{(d)}$  be the extensions of  $\hat{u}^{(1)}, \dots, \hat{u}^{(d)}$  as defined by equation (B.2). We then have, for any  $s = 1, 2, \dots, d$ ,

$$\begin{aligned} \langle \hat{v}^{(s)}, \sqrt{\lambda_r} \psi_r \rangle_{\mathcal{H}} &= \left\langle \frac{1}{\sqrt{\hat{\lambda}_s n}} \sum_{i=1}^n \kappa(\cdot, X_i) \hat{u}_i^{(s)}, \sqrt{\lambda_r} \psi_r \right\rangle_{\mathcal{H}} \\ &= \left\langle \frac{1}{\sqrt{\hat{\lambda}_s n}} \sum_{i=1}^n \sum_{r'} \sqrt{\lambda_{r'}} \psi_{r'}(X_i) \sqrt{\lambda_{r'}} \psi_{r'} \hat{u}_i^{(s)}, \sqrt{\lambda_r} \psi_r \right\rangle_{\mathcal{H}} \\ &= \frac{1}{\sqrt{\hat{\lambda}_s n}} \sum_{i=1}^n \psi_r(X_i) \sqrt{\lambda_r} \hat{u}_i^{(s)} \\ &= \frac{1}{\sqrt{\hat{\lambda}_s n}} \langle \hat{u}^{(s)}, \Psi_{r,n} \rangle_{\mathbb{R}^n}. \end{aligned}$$

We thus have

$$(A.1) \quad \hat{u}_i^{(s)} (\hat{u}^{(s)})^T \Psi_{r,n} = \hat{u}_i^{(s)} \langle \hat{u}^{(s)}, \Psi_{r,n} \rangle_{\mathbb{R}^n} = \hat{v}^{(s)}(X_i) \langle \hat{v}^{(s)}, \sqrt{\lambda_r} \psi_r \rangle_{\mathcal{H}}.$$

Now let  $\xi^{(s)}(X) = \sum_{r=1}^{\infty} \langle \hat{v}^{(s)}, \psi_r \sqrt{\lambda_r} \rangle_{\mathcal{H}} \hat{v}^{(s)}(X) \sqrt{\lambda_r} \psi_r \in \mathcal{H}$ .  $\xi^{(s)}(X)$  is the embedding of the sequence  $(\langle \hat{v}^{(s)}, \sqrt{\lambda_r} \psi_r \rangle_{\mathcal{H}} \hat{v}^{(s)}(X))_{r=1}^{\infty} \in l_2$  into  $\mathcal{H}$ ; see equation (2.2). By equation (A.1) and the definition of  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  [equation (2.3)], the  $ij$ th entry of  $\mathcal{P}_{\mathbf{K}}\mathbf{K}$  can be written as

$$\sum_{s=1}^d \sum_{r=1}^{\infty} \hat{u}_i^{(s)} (\hat{u}^{(s)})^T \Psi_{r,n} \Psi_{r,n}^T \hat{u}^{(s)} \hat{u}_j^{(s)} = \sum_{s=1}^d \langle \xi^{(s)}(X_i), \xi^{(s)}(X_j) \rangle_{\mathcal{H}}.$$

We note that, by the reproducing kernel property of  $\kappa(\cdot, x)$ ,

$$\begin{aligned} \xi^{(s)}(X) &= \sum_{r=1}^{\infty} \langle \hat{v}^{(s)}, \psi_r \sqrt{\lambda_r} \rangle_{\mathcal{H}} \hat{v}^{(s)}(X) \sqrt{\lambda_r} \psi_r \\ &= \sum_{r=1}^{\infty} \langle \hat{v}^{(s)}, \psi_r \sqrt{\lambda_r} \rangle_{\mathcal{H}} \langle \hat{v}^{(s)}, \kappa(\cdot, X) \rangle_{\mathcal{H}} \sqrt{\lambda_r} \psi_r \end{aligned}$$

$$\begin{aligned} &= \langle \hat{v}^{(s)}, \kappa(\cdot, X) \rangle_{\mathcal{H}} \sum_{r=1}^{\infty} \langle \hat{v}^{(s)}, \psi_r \sqrt{\lambda_r} \rangle_{\mathcal{H}} \sqrt{\lambda_r} \psi_r \\ &= \langle \hat{v}^{(s)}, \kappa(\cdot, X) \rangle_{\mathcal{H}} \hat{v}^{(s)}. \end{aligned}$$

As the  $\hat{v}^{(s)}$  are orthogonal with respect to  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ , the  $ij$ th entry of  $\mathcal{P}_{\mathbf{K}}\mathbf{K}$  can also be written as

$$\begin{aligned} &\sum_{s=1}^d \langle \xi^{(s)}(X_i), \xi^{(s)}(X_j) \rangle_{\mathcal{H}} \\ &= \sum_{s=1}^d \langle \hat{v}^{(s)}, \kappa(\cdot, X_i) \rangle_{\mathcal{H}} \langle \hat{v}^{(s)}, \hat{v}^{(s)} \rangle_{\mathcal{H}} \langle \hat{v}^{(s)}, \kappa(\cdot, X_j) \rangle_{\mathcal{H}} \\ &= \sum_{s=1}^d \sum_{s'=1}^d \langle \hat{v}^{(s)}, \kappa(\cdot, X_i) \rangle_{\mathcal{H}} \langle \hat{v}^{(s)}, \hat{v}^{(s')} \rangle_{\mathcal{H}} \langle \hat{v}^{(s')}, \kappa(\cdot, X_j) \rangle_{\mathcal{H}} \\ &= \left\langle \sum_{s=1}^d \langle \hat{v}^{(s)}, \kappa(\cdot, X_i) \rangle_{\mathcal{H}} \hat{v}^{(s)}, \sum_{s=1}^d \langle \hat{v}^{(s)}, \kappa(\cdot, X_j) \rangle_{\mathcal{H}} \hat{v}^{(s)} \right\rangle_{\mathcal{H}} \\ &= \langle \hat{\mathcal{P}}_d \kappa(\cdot, X_i), \hat{\mathcal{P}}_d \kappa(\cdot, X_j) \rangle_{\mathcal{H}}. \end{aligned}$$

As the  $\hat{\mathcal{P}}_d \kappa(\cdot, \cdot)$  lies in a  $d$ -dimensional subspace of  $\mathcal{H}$ , they can be isometrically embedded into  $\mathbb{R}^d$ . Thus there exists a matrix  $\mathbf{X} \in \mathcal{M}_{n,d}(\mathbb{R})$  such that  $\mathbf{X}\mathbf{X}^T = \mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}\mathbf{U}_{\mathbf{K}}^T$  and that the rows of  $\mathbf{X}$  correspond to the projections  $\hat{\mathcal{P}}_d \kappa(\cdot, X_i)$ . Therefore, there exists a unitary matrix  $\mathbf{W} \in \mathcal{M}_d(\mathbf{R})$  such that  $\mathbf{X} = \mathbf{U}_{\mathbf{K}}\mathbf{S}_{\mathbf{K}}^{1/2}\mathbf{W}$  as desired.  $\square$

LEMMA A.1. *Let  $\mathbf{A}$  and  $\mathbf{B}$  be  $n \times n$  positive semidefinite matrices with  $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{B}) = d$ . Let  $\mathbf{X}, \mathbf{Y} \in \mathcal{M}_{n,d}(\mathbb{R})$  be of full column rank such that  $\mathbf{X}\mathbf{X}^T = \mathbf{A}$  and  $\mathbf{Y}\mathbf{Y}^T = \mathbf{B}$ . Let  $\delta$  be the smallest nonzero eigenvalue of  $\mathbf{B}$ . Then there exists an orthogonal matrix  $\mathbf{W} \in \mathcal{M}_d(\mathbb{R})$  such that*

$$(A.2) \quad \|\mathbf{X}\mathbf{W} - \mathbf{Y}\|_F \leq \frac{\|\mathbf{A} - \mathbf{B}\|(\sqrt{d\|\mathbf{A}\|} + \sqrt{d\|\mathbf{B}\|})}{\delta}.$$

PROOF. Let  $\mathbf{R} = \mathbf{A} - \mathbf{B}$ . As  $\mathbf{Y}$  is of full column rank,  $\mathbf{Y}^T\mathbf{Y}$  is invertible, and its smallest eigenvalue is  $\delta$ . We then have

$$\mathbf{Y} = \mathbf{X}\mathbf{X}^T\mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1} - \mathbf{R}\mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1}.$$

Let  $\mathbf{T} = \mathbf{X}^T\mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1}$ . We then have

$$\mathbf{T}^T\mathbf{T} - \mathbf{I} = (\mathbf{Y}^T\mathbf{Y})^{-1}\mathbf{Y}^T\mathbf{X}\mathbf{X}^T\mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1} - \mathbf{I} = (\mathbf{Y}^T\mathbf{Y})^{-1}\mathbf{Y}^T\mathbf{R}\mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1}.$$

Therefore,

$$-(\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y}^T \|\mathbf{R}\| \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1} \leq \mathbf{T}^T \mathbf{T} - \mathbf{I} \leq (\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y}^T \|\mathbf{R}\| \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1},$$

where  $\leq$  refers to the positive semi-definite ordering for matrices. We thus have

$$\|\mathbf{T}^T \mathbf{T} - \mathbf{I}\|_F \leq \|\mathbf{R}\| \cdot \|(\mathbf{Y}^T \mathbf{Y})^{-1}\|_F \leq \sqrt{d} \|\mathbf{R}\| \cdot \|(\mathbf{Y}^T \mathbf{Y})^{-1}\| \leq \frac{\|\mathbf{R}\| \sqrt{d}}{\delta}.$$

Now let  $\mathbf{W}$  be the orthogonal matrix in the polar decomposition  $\mathbf{T} = \mathbf{W}(\mathbf{T}^T \mathbf{T})^{1/2}$ . We then have

$$\begin{aligned} \|\mathbf{XW} - \mathbf{Y}\|_F &\leq \|\mathbf{XW} - \mathbf{XT}\|_F + \|\mathbf{XT} - \mathbf{Y}\|_F \\ &\leq \|\mathbf{X}\| \cdot \|(\mathbf{T}^T \mathbf{T})^{1/2} - \mathbf{I}\|_F + \|\mathbf{R}\| \cdot \|\mathbf{Y}(\mathbf{Y}^T \mathbf{Y})^{-1}\|_F \\ &\leq \|\mathbf{X}\| \cdot \|(\mathbf{T}^T \mathbf{T})^{1/2} - \mathbf{I}\|_F + \|\mathbf{R}\| \cdot \|\mathbf{Y}\| \cdot \|(\mathbf{Y}^T \mathbf{Y})^{-1}\|_F. \end{aligned}$$

Now,  $\|(\mathbf{T}^T \mathbf{T})^{1/2} - \mathbf{I}\|_F \leq \|\mathbf{T}^T \mathbf{T} - \mathbf{I}\|_F$ . Indeed,

$$\begin{aligned} \|(\mathbf{T}^T \mathbf{T})^{1/2} - \mathbf{I}\|_F^2 &= \sum_{i=1}^d (\lambda_i(\mathbf{T}^T \mathbf{T})^{1/2} - 1)^2 \leq \sum_{i=1}^d (\lambda_i(\mathbf{T}^T \mathbf{T}) - 1)^2 \\ &= \|\mathbf{T}^T \mathbf{T} - \mathbf{I}\|_F^2. \end{aligned}$$

We thus have

$$\|\mathbf{XW} - \mathbf{Y}\| \leq (\|\mathbf{X}\| + \|\mathbf{Y}\|) \frac{\|\mathbf{R}\| \sqrt{d}}{\delta},$$

and equation (A.2) follows.  $\square$

**PROPOSITION A.2.** *Let  $\kappa$  be a universal kernel on  $\mathcal{X}$ , and let  $\Phi : \mathcal{X} \mapsto l_2$  be a feature map of  $\kappa$ . Let  $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \dots$  be the sequence of classifiers of the form in equation (4.1). Then*

$$(A.3) \quad \lim_{d \rightarrow \infty} \inf_{f \in \mathcal{C}^{(d)}} R_\varphi(f) = R_\varphi^*.$$

**PROOF.** We note that this result is a slight variation of Lemma 1 in [24]. For completeness, we sketch its proof here. Let  $f^*$  be the function defined by

$$f^*(x) = \inf_{\alpha \in \mathbb{R}} \{ \eta(x) \varphi(\alpha) + (1 - \eta(x)) \varphi(-\alpha) \},$$

where  $\eta(x) = \mathbb{P}[Y = 1 | X = x]$ . Then  $R_\varphi^* = \mathbb{E}[f^*]$ . Now, for a given  $\beta \in [0, 1/2]$ , let  $H_\beta = \{x : |\eta(x) - 1/2| > \beta\}$ , and let  $\bar{H}_\beta$  be the complement of  $H_\beta$ . We consider the decomposition

$$R_\varphi^* = \mathbb{E}[f^*(X) \mathbf{1}\{X \in H_\beta\}] + \mathbb{E}[f^*(X) \mathbf{1}\{X \in \bar{H}_\beta\}].$$

The restriction of  $f^*$  to  $\bar{H}_\beta$  is measurable with range  $[-C_\beta, C_\beta]$  for some finite constant  $C_\beta > 0$ . The set of functions  $\langle w, \Phi \rangle_{\mathcal{H}}$  is dense in  $\mathcal{C}(X)$  and hence also dense in  $L^1(\mathcal{X}, F_{\mathcal{X}})$ . Thus, for any  $\varepsilon > 0$ , there exists a  $w \in \mathcal{H}$  such that

$$\mathbb{E}[f^*(X)\mathbf{1}\{X \in \bar{H}_\beta\}] - \mathbb{E}[\langle w, \Phi(X) \rangle_{\mathcal{H}}\mathbf{1}\{X \in \bar{H}_\beta\}] < \varepsilon.$$

Furthermore,  $\mathbb{E}[f^*(X)\mathbf{1}\{X \in H_\beta\}] \rightarrow 0$  as  $\beta \rightarrow 1/2$ . Indeed,  $H_{1/2} = \{x : \eta(x) \in \{0, 1\}\}$  so we can select  $\alpha$  so that  $\varphi(\alpha) = 0$  if  $\eta(x) = 1$  and  $\varphi(-\alpha) = 0$  if  $\eta(x) = 0$ . To complete the proof, we note that the  $\mathcal{C}^{(d)}$  are nested, that is,  $\mathcal{C}^{(d)} \subset \mathcal{C}^{(d+1)}$ . Hence  $\inf_{f \in \mathcal{C}^{(d)}} R_\varphi(f)$  is a decreasing sequence that converges to  $R_\varphi^*$  as desired.  $\square$

### APPENDIX B: SPECTRA OF INTEGRAL OPERATORS AND KERNEL MATRICES

We can tie the spectrum and eigenvectors of  $\mathbf{K}$  to the spectrum and eigenfunctions of  $\mathcal{H}$  by constructing an extension operator  $\mathcal{K}_{\mathcal{H},n}$  for  $\mathbf{K}$  and relating the spectra of  $\mathcal{H}$  to that of  $\mathcal{K}_{\mathcal{H},n}$  [28]. Let  $\mathcal{H}$  be the reproducing kernel Hilbert space for  $\kappa$ . Let  $\mathcal{K}_{\mathcal{H}} : \mathcal{H} \mapsto \mathcal{H}$  and  $\mathcal{K}_{\mathcal{H},n} : \mathcal{H} \mapsto \mathcal{H}$  be the linear operators defined by

$$\begin{aligned} \mathcal{K}_{\mathcal{H}}\eta &= \int_{\mathcal{X}} \langle \eta, \kappa(\cdot, x) \rangle_{\mathcal{H}} \kappa(\cdot, x) dF(x), \\ \mathcal{K}_{\mathcal{H},n}\eta &= \frac{1}{n} \sum_{i=1}^n \langle \eta, \kappa(\cdot, X_i) \rangle_{\mathcal{H}} \kappa(\cdot, X_i). \end{aligned}$$

The operators  $\mathcal{K}_{\mathcal{H}}$  and  $\mathcal{K}_{\mathcal{H},n}$  are defined on the same Hilbert space  $\mathcal{H}$ , in contrast to  $\mathcal{H}$  and  $\mathbf{K}$  which are defined on the different spaces  $L^2(\mathcal{X}, F)$  and  $\mathbb{R}^n$ , respectively. Thus, we can relate the spectra of  $\mathcal{K}_{\mathcal{H}}$  and  $\mathcal{K}_{\mathcal{H},n}$ . Furthermore, we can also relate the spectra of  $\mathcal{H}$  and  $\mathcal{K}_{\mathcal{H}}$  as well as the spectra of  $\mathbf{K}$  and  $\mathcal{K}_{\mathcal{H},n}$ , therefore giving us a relationship between the spectra of  $\mathcal{H}$  and  $\mathbf{K}$ . A precise statement of the relationships is contained in the following results.

**PROPOSITION B.1** ([28, 38]). *The operators  $\mathcal{K}_{\mathcal{H}}$  and  $\mathcal{K}_{\mathcal{H},n}$  are positive, self-adjoint operators and are of trace class with  $\mathcal{K}_{\mathcal{H},n}$  being of finite rank. The spectra of  $\mathcal{H}$  and  $\mathcal{K}_{\mathcal{H}}$  are contained in  $[0, 1]$  and are the same, possibly up to the zero eigenvalues. If  $\lambda$  is a nonzero eigenvalue of  $\mathcal{H}$  and  $u$  and  $v$  are associated eigenfunction of  $\mathcal{H}$  and  $\mathcal{K}_{\mathcal{H}}$ , normalized to norm 1 in  $L^2(\mathcal{X}, F)$  and  $\mathcal{H}$ , respectively, then*

$$\begin{aligned} (B.1) \quad u(x) &= \frac{v(x)}{\sqrt{\lambda}} \quad \text{for } x \in \text{supp}(F); \\ v(x) &= \frac{1}{\sqrt{\lambda}} \int_{\mathcal{X}} \kappa(x, x')u(x') dF(x'). \end{aligned}$$

Similarly, the spectra of  $\mathbf{K}/n$  and  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$  are contained in  $[0, 1]$  and are the same, possibly up to the zero eigenvalues. If  $\hat{\lambda}$  is a nonzero eigenvalue of  $\mathbf{K}$  and  $\hat{u}$  and  $\hat{v}$  are the corresponding eigenvector and eigenfunction of  $\mathbf{K}/n$  and  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$ , normalized to norm 1 in  $\mathbb{R}^n$  and  $\mathcal{H}$ , respectively, then

$$(B.2) \quad \hat{u}_i = \frac{\hat{v}(x_i)}{\sqrt{\hat{\lambda}}}; \quad \hat{v}(\cdot) = \frac{1}{\sqrt{\hat{\lambda}n}} \sum_{i=1}^n \kappa(\cdot, x_i) \hat{u}_i.$$

Equation (B.2) in Proposition B.1 states that an eigenvector  $\hat{u}$  of  $\mathbf{K}/n$ , which is only defined for  $X_1, X_2, \dots, X_n$ , can be extended to an eigenfunction  $\hat{v} \in \mathcal{H}$  of  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$  defined for all  $x \in \mathcal{X}$ , and furthermore, that  $\hat{u}_i = \hat{v}(X_i)$  for all  $i = 1, 2, \dots, n$ .

**THEOREM B.2** ([28, 41]). *Let  $\tau > 0$  be arbitrary. Then with probability at least  $1 - 2e^{-\tau}$ ,*

$$(B.3) \quad \|\mathcal{K}\mathcal{H} - \mathcal{K}\mathcal{H}_{\mathcal{H},n}\|_{\text{HS}} \leq 2\sqrt{2}\sqrt{\frac{\tau}{n}},$$

where  $\|\cdot\|_{\text{HS}}$  is the Hilbert–Schmidt norm. Let  $\{\lambda_j\}$  be a decreasing enumeration of the eigenvalues for  $\mathcal{K}\mathcal{H}$ , and let  $\{\hat{\lambda}_j\}$  be an extended decreasing enumeration of  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$ ; that is,  $\hat{\lambda}_j$  is either an eigenvalue of  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$  or  $\hat{\lambda}_j = 0$ . Then the above bound and a Lidskii theorem for infinite-dimensional operators [22] yields

$$(B.4) \quad \left(\sum_{j \geq 1} (\lambda_j - \hat{\lambda}_j)^2\right)^{1/2} \leq 2\sqrt{2}\sqrt{\frac{\tau}{n}}$$

with probability at least  $1 - 2e^{-\tau}$ . For a given  $d \geq 1$  and  $\tau > 0$ , if the number  $n$  of samples  $X_i \sim F$  satisfies

$$4\sqrt{2}\sqrt{\frac{\tau}{n}} < \lambda_d - \lambda_{d+1},$$

then with probability greater than  $1 - 2e^{-\tau}$ ,

$$(B.5) \quad \|\mathcal{P}_d - \hat{\mathcal{P}}_d\|_{\text{HS}} \leq \frac{2\sqrt{2}\sqrt{\tau}}{(\lambda_d - \lambda_{d+1})\sqrt{n}},$$

where  $\mathcal{P}_d$  is the projection onto the subspace spanned by the eigenfunctions corresponding to the  $d$  largest eigenvalues of  $\mathcal{K}$ , and  $\hat{\mathcal{P}}_d$  is the projection onto the subspace spanned by the eigenfunctions corresponding to the  $d$  largest eigenvalues of  $\mathcal{K}\mathcal{H}_{\mathcal{H},n}$ .

**Acknowledgements.** The authors thank the anonymous referees for their comments which have improved the presentation and quality of the paper.

REFERENCES

[1] AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.

- [2] BARTLETT, P. L., JORDAN, M. I. and MCAULIFFE, J. D. (2006). Convexity, classification, and risk bounds. *J. Amer. Statist. Assoc.* **101** 138–156. [MR2268032](#)
- [3] BELKIN, M. and NIYOGI, P. (2005). Towards a theoretical foundation for Laplacian-based manifold methods. In *Proceedings of the 18th Conference on Learning Theory (COLT 2005)* 486–500. Springer, Berlin.
- [4] BENGIO, Y., VINCENT, P., PAIEMENT, J. F., DELALLEAU, O., OUMET, M. and ROUX, N. L. (2004). Learning eigenfunctions links spectral embedding and kernel PCA. *Neural Comput.* **16** 2197–2219.
- [5] BIAU, G., BUNEA, F. and WEGKAMP, M. H. (2005). Functional classification in Hilbert spaces. *IEEE Trans. Inform. Theory* **51** 2163–2172. [MR2235289](#)
- [6] BLANCHARD, G., BOUSQUET, O. and MASSART, P. (2008). Statistical performance of support vector machines. *Ann. Statist.* **36** 489–531. [MR2396805](#)
- [7] BLANCHARD, G., BOUSQUET, O. and ZWALD, L. (2007). Statistical properties of kernel principal component analysis. *Machine Learning* **66** 259–294.
- [8] BOLLOBÁS, B., JANSON, S. and RIORDAN, O. (2007). The phase transition in inhomogeneous random graphs. *Random Structures Algorithms* **31** 3–122. [MR2337396](#)
- [9] BRAUN, M. L. (2006). Accurate error bounds for the eigenvalues of the kernel matrix. *J. Mach. Learn. Res.* **7** 2303–2328. [MR2274441](#)
- [10] CHATTERJEE, S. (2012). Matrix estimation by universal singular value thresholding. Preprint. Available at <http://arxiv.org/abs/1212.1247>.
- [11] CHAUDHURI, K., CHUNG, F. and TSIATAS, A. (2012). Spectral partitioning of graphs with general degrees and the extended planted partition model. In *Proceedings of the 25th Conference on Learning Theory*. Springer, Berlin.
- [12] CUCKER, F. and SMALE, S. (2002). On the mathematical foundations of learning. *Bull. Amer. Math. Soc. (N.S.)* **39** 1–49 (electronic). [MR1864085](#)
- [13] DAVIS, C. and KAHAN, W. M. (1970). The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* **7** 1–46. [MR0264450](#)
- [14] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. (1996). *A Probabilistic Theory of Pattern Recognition. Applications of Mathematics (New York)* **31**. Springer, New York. [MR1383093](#)
- [15] DIACONIS, P. and JANSON, S. (2008). Graph limits and exchangeable random graphs. *Rend. Mat. Appl. (7)* **28** 33–61. [MR2463439](#)
- [16] FISHKIND, D. E., SUSSMAN, D. L., TANG, M., VOGELSTEIN, J. T. and PRIEBE, C. E. (2013). Consistent adjacency-spectral partitioning for the stochastic block model when the model parameters are unknown. *SIAM J. Matrix Anal. Appl.* **34** 23–39. [MR3032990](#)
- [17] HEIN, M., AUDIBERT, J.-Y. and VON LUXBURG, U. (2005). From graphs to manifolds—Weak and strong pointwise consistency of graph Laplacians. In *Proceedings of the 18th Conference on Learning Theory (COLT 2005)* 470–485. Springer, Berlin.
- [18] HEIN, M., AUDIBERT, J. Y. and VON LUXBURG, U. (2007). Convergence of graph Laplacians on random neighbourhood graphs. *J. Mach. Learn. Res.* **8** 1325–1370.
- [19] HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* **97** 1090–1098. [MR1951262](#)
- [20] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- [21] KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* (3) **83** 016107, 10. [MR2788206](#)
- [22] KATO, T. (1987). Variation of discrete spectra. *Comm. Math. Phys.* **111** 501–504. [MR0900507](#)
- [23] KOLTCHINSKII, V. and GINÉ, E. (2000). Random matrix approximation of spectra of integral operators. *Bernoulli* **6** 113–167. [MR1781185](#)
- [24] LUGOSI, G. and VAYATIS, N. (2004). On the Bayes-risk consistency of regularized boosting methods. *Ann. Statist.* **32** 30–55. [MR2051000](#)

- [25] MICCHELLI, C. A., XU, Y. and ZHANG, H. (2006). Universal kernels. *J. Mach. Learn. Res.* **7** 2651–2667. [MR2274454](#)
- [26] OLIVEIRA, R. I. (2010). Concentration of the adjacency matrix and of the Laplacian in random graphs with independent edges. Preprint. Available at <http://arxiv.org/abs/0911.0600>.
- [27] ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. [MR2893856](#)
- [28] ROSASCO, L., BELKIN, M. and DE VITO, E. (2010). On learning with integral operators. *J. Mach. Learn. Res.* **11** 905–934. [MR2600634](#)
- [29] SHAWE-TAYLOR, J., WILLIAMS, C. K. I., CRISTIANINI, N. and KANDOLA, J. (2005). On the eigenspectrum of the Gram matrix and the generalization error of kernel-PCA. *IEEE Trans. Inform. Theory* **51** 2510–2522. [MR2246374](#)
- [30] SRIPERUMBUDUR, B. K., FUKUMIZU, K. and LANCKRIET, G. R. G. (2011). Universality, characteristic kernels and RKHS embedding of measures. *J. Mach. Learn. Res.* **12** 2389–2410. [MR2825431](#)
- [31] STEINWART, I. (2002). On the influence of the kernel on the consistency of support vector machines. *J. Mach. Learn. Res.* **2** 67–93. [MR1883281](#)
- [32] STEINWART, I. (2002). Support vector machines are universally consistent. *J. Complexity* **18** 768–791. [MR1928806](#)
- [33] STEINWART, I. and SCOVEL, C. (2012). Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constr. Approx.* **35** 363–417. [MR2914365](#)
- [34] STONE, C. J. (1977). Consistent nonparametric regression. *Ann. Statist.* **5** 595–620.
- [35] SUSSMAN, D. L., TANG, M., FISHKIND, D. E. and PRIEBE, C. E. (2012). A consistent adjacency spectral embedding for stochastic blockmodel graphs. *J. Amer. Statist. Assoc.* **107** 1119–1128. [MR3010899](#)
- [36] SUSSMAN, D. L., TANG, M. and PRIEBE, C. E. (2012). Consistent latent position estimation and vertex classification for random dot product graphs. Preprint. Available at <http://arxiv.org/abs/1207.6745>.
- [37] TROPP, J. A. (2011). Freedman’s inequality for matrix martingales. *Electron. Commun. Probab.* **16** 262–270. [MR2802042](#)
- [38] VON LUXBURG, U., BELKIN, M. and BOUSQUET, O. (2008). Consistency of spectral clustering. *Ann. Statist.* **36** 555–586. [MR2396807](#)
- [39] YOUNG, S. J. and SCHEINERMAN, E. R. (2007). Random dot product graph models for social networks. In *Algorithms and Models for the Web-Graph. Lecture Notes in Computer Science* **4863** 138–149. Springer, Berlin. [MR2504912](#)
- [40] ZHANG, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. *Ann. Statist.* **32** 56–85. [MR2051001](#)
- [41] ZWALD, L. and BLANCHARD, G. (2006). On the convergence of eigenspaces in kernel principal components analysis. In *Advances in Neural Information Processing Systems (NIPS 05)* **18** 1649–1656. MIT Press, Cambridge.

DEPARTMENT OF APPLIED MATHEMATICS  
AND STATISTICS  
JOHNS HOPKINS UNIVERSITY  
3400 N. CHARLES ST  
BALTIMORE, MARYLAND 21218  
USA  
E-MAIL: [mtang10@jhu.edu](mailto:mtang10@jhu.edu)  
[dsussma3@jhu.edu](mailto:dsussma3@jhu.edu)  
[cep@jhu.edu](mailto:cep@jhu.edu)