

Kernel ridge vs. principal component regression: Minimax bounds and the qualification of regularization operators

Lee H. Dicker

*Department of Statistics and Biostatistics
Rutgers University
Piscataway, NJ
e-mail: ldicker@stat.rutgers.edu*

Dean P. Foster

*Amazon Inc.
New York, NY
e-mail: dean@foster.net*

and

Daniel Hsu

*Department of Computer Science
Columbia University
New York, NY
e-mail: djhsu@cs.columbia.edu*

Abstract: Regularization is an essential element of virtually all kernel methods for nonparametric regression problems. A critical factor in the effectiveness of a given kernel method is the type of regularization that is employed. This article compares and contrasts members from a general class of regularization techniques, which notably includes ridge regression and principal component regression. We derive an explicit finite-sample risk bound for regularization-based estimators that simultaneously accounts for (i) the structure of the ambient function space, (ii) the regularity of the true regression function, and (iii) the adaptability (or *qualification*) of the regularization. A simple consequence of this upper bound is that the risk of the regularization-based estimators matches the minimax rate in a variety of settings. The general bound also illustrates how some regularization techniques are more adaptable than others to favorable regularity properties that the true regression function may possess. This, in particular, demonstrates a striking difference between kernel ridge regression and kernel principal component regression. Our theoretical results are supported by numerical experiments.

MSC 2010 subject classifications: 62G08.

*DH acknowledges support from NSF grants DMR-1534910 and IIS-1563785 and a Sloan Research Fellowship; LHD's work was partially supported by NSF grants DMS-1208785 and DMS-1454817.

Keywords and phrases: Nonparametric regression, reproducing kernel Hilbert space.

Received August 2016.

1. Introduction

Suppose that the observed data consists of $\mathbf{z}_i = (y_i, \mathbf{x}_i)$, $i = 1, \dots, n$, where $y_i \in \mathcal{Y} \subseteq \mathbb{R}$ and $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^d$. Suppose further that $\mathbf{z}_1, \dots, \mathbf{z}_n \sim \rho$ are iid from some probability distribution ρ on $\mathcal{Y} \times \mathcal{X}$. Let $\rho(\cdot | \mathbf{x})$ denote the conditional distribution of y_i given $\mathbf{x}_i = \mathbf{x} \in \mathcal{X}$ and let $\rho_{\mathcal{X}}$ denote the marginal distribution of \mathbf{x}_i . Our goal is to use the available data to estimate the regression function of y on \mathbf{x} ,

$$f^\dagger(\mathbf{x}) = \int_{\mathcal{Y}} y \, d\rho(y | \mathbf{x}),$$

which minimizes the mean-squared prediction error

$$\int_{\mathcal{Y} \times \mathcal{X}} (y - f(\mathbf{x}))^2 \, d\rho(y, \mathbf{x})$$

over $\rho_{\mathcal{X}}$ -measurable functions $f: \mathcal{X} \rightarrow \mathbb{R}$. More specifically, for an estimator \hat{f} define the risk

$$\mathcal{R}_\rho(\hat{f}) = \mathbb{E} \left[\int (f^\dagger(\mathbf{x}) - \hat{f}(\mathbf{x}))^2 \, d\rho_{\mathcal{X}}(\mathbf{x}) \right] = \mathbb{E} \left[\|f^\dagger - \hat{f}\|_{\rho_{\mathcal{X}}}^2 \right], \quad (1)$$

where the expectation is computed over $\mathbf{z}_1, \dots, \mathbf{z}_n$, and $\|\cdot\|_{\rho_{\mathcal{X}}}$ denotes the norm on $L^2(\rho_{\mathcal{X}})$; we seek estimators \hat{f} which minimize $\mathcal{R}_\rho(\hat{f})$.

This is a version of the random design nonparametric regression problem. There is a vast literature on nonparametric regression, along with a huge variety of corresponding methods (e.g., Györfi et al., 2002; Wasserman, 2006). In this paper, we focus on regularization and kernel methods for estimating f^\dagger .¹ Most of our results apply to general regularization operators. However, our motivating examples are two well-known regularization techniques: Kernel ridge regression (which we refer to as “KRR”; KRR is also known as Tikhonov regularization) and kernel principal component regression (“KPCR”; also known as spectral cut-off regularization).

Our main theorem is a new upper bound on the risk of a general class of kernel-regularization methods, which includes both KRR and KPCR (Theorem 1). The theorem substantially generalizes previously published bounds (see Section 2 for a discussion of related work) and illustrates the dependence of the risk on three important features: (i) the structure of the ambient reproducing kernel Hilbert space (RKHS), (ii) the specific regularization technique employed, and (iii) the regularity (often interpreted as smoothness) of the function to be

¹Here, we mean “kernel” as in *reproducing kernel Hilbert space*, rather than *kernel-smoothing*, which is another popular approach to nonparametric regression.

estimated. One consequence of the theorem is that the regularization methods studied in this paper (including KRR and KPCR) achieve the minimax rate for estimating f^\dagger in a variety of settings. A second consequence is that certain regularization methods (including KPCR, but not KRR) may adapt to favorable regularity of f^\dagger to attain even faster convergence rates, while others (notably KRR) are limited in this regard due to a well-known *saturation* effect (Neubauer, 1997; Mathé, 2005; Bauer et al., 2007). This illustrates a striking advantage that KPCR may have over KRR in these settings.

2. Related work

Kernel ridge regression has been studied extensively in the literature. Indeed, bounds for KRR that are similar to our Theorem 1 have been derived by Caponnetto and De Vito (2007). Moreover, it is well-known that KRR is minimax in many of the settings considered in this paper, such as those described in Corollaries 1–3 (Caponnetto and De Vito, 2007; Zhang, 2005; Zhang et al., 2013; Lin et al., 2016). However, these cited results apply only to KRR, while the results in this paper apply to a substantially larger class of regularization operators (including, for example, KPCR).

Beyond KRR, there has also been significant research into more general regularization methods, like those considered in this paper. However, our bounds are sharper than previously published results on general regularization operators. For instance, unlike our Theorem 1, the bounds of Bauer et al. (2007) do not illustrate the dependence of the risk on the ambient Hilbert space. Thus, while our approach immediately implies that many of the regularization methods under consideration are minimax optimal, it seems difficult (if not impossible) to draw this conclusion using the approach of Bauer et al. (2007). General regularization operators are also studied by Caponnetto and Yao (2010), but their results apply to semi-supervised settings where an additional pool of unlabeled data is available.

One of the major practical implications of this paper is that KPCR may have significant advantages over KRR in some settings. This has been observed previously by other researchers; others have even noted that Tikhonov regularization (KRR) saturates, while spectral cut-off regularization (KPCR) does not (Mathé, 2005; Bauer et al., 2007; Lo Gerfo et al., 2008). Our results (Theorem 1 and Corollaries 1–3) sharpen these observations by precisely quantifying the advantages of unsaturated regularization operators in terms of adaptability and minimaxity. In other related work, Dhillon et al. (2013) have illustrated the potential advantages of KPCR over KRR in finite-dimensional problems with linear kernels; though their work is not framed in terms of saturation and general regularization operators, it relies on similar concepts.

After the initial submission of our work, we became aware of simultaneous independent work by Blanchard and Mücke (2016) that proves upper bounds on the risk for a comparable class of regularization methods (which includes KRR and KPCR), as well as minimax lower bounds that match the upper

bounds; their results are specialized to RKHS's where the covariance operator has polynomially decaying eigenvalues. Compared to that work, our results require a weaker moment condition on the noise for risk bounds, apply to a much broader class of RKHS's, and also consider target functions that live in finite-dimensional subspaces (see Proposition 1). Another distinguishing feature of the present work is that it contains numerical experiments that illustrate the implications of our theoretical bounds in some practical settings (Section 6).

Other important recent work was pointed out to us by a reviewer. Optimal rates in expectation for KRR and more general regularization methods are derived by Lin et al. (2016) and Guo et al. (2016), respectively, for distributed learning problems. The bounds in the latter paper, when specialized to the non-distributed setting, are similar to those in this paper and by Blanchard and Mücke (2016). However, the results in this paper are more general (for non-distributed learning), because they do not require continuity of the kernel and have weaker conditions on the output variable.

The main engine behind the technical results in this paper is a collection of large-deviation results for Hilbert-Schmidt operators. The required machinery is developed in the Appendices. These results build on straightforward extensions of results of Tropp (2015). Our most precise results for KPCR and achieving parametric rates for estimation over finite-dimensional subspaces (Proposition 1) rely on slightly different arguments, which are based on well-known eigenvalue perturbation results that have been adapted to handle Hilbert-Schmidt operators (e.g., the Davis-Kahan $\sin \Theta$ theorem (Davis and Kahan, 1970)).

3. Statistical setting and assumptions

Our basic assumption on the distribution of $\mathbf{z} = (y, \mathbf{x}) \sim \rho$ is that the residual variance is bounded; more specifically, we assume that there exists a constant $\sigma^2 > 0$ such that

$$\int_{\mathcal{Y}} \left[y - f^\dagger(\mathbf{x}) \right]^2 d\rho(y | \mathbf{x}) \leq \sigma^2 \quad (2)$$

for almost all $\mathbf{x} \in \mathcal{X}$. Zhang et al. (2013) also assume (2); this assumption is slightly weaker than the analogous assumption used by Bauer et al. (2007) (equation (1) in their paper). Note that (2) holds if y is bounded almost surely or if $y = f^\dagger(\mathbf{x}) + \epsilon$, where ϵ is independent of \mathbf{x} , $\mathbb{E}(\epsilon) = 0$, and $\text{Var}(\epsilon) = \sigma^2$.

Let $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a symmetric positive-definite kernel function. We assume that K is bounded—i.e., that there exists $\kappa^2 > 0$ such that $\sup_{\mathbf{x} \in \mathcal{X}} K(\mathbf{x}, \mathbf{x}) \leq \kappa^2$. Additionally, we assume that there is a countable orthonormal basis $\{\psi_j\}_{j=1}^\infty \subseteq L^2(\rho_{\mathcal{X}})$ and a sequence of corresponding real numbers $t_1^2 \geq t_2^2 \geq \dots \geq 0$ such that

$$K(\mathbf{x}, \tilde{\mathbf{x}}) = \sum_{j=1}^{\infty} t_j^2 \psi_j(\mathbf{x}) \psi_j(\tilde{\mathbf{x}}), \quad \mathbf{x}, \tilde{\mathbf{x}} \in \mathcal{X}, \quad (3)$$

where the convergence is absolute and uniform over $\mathbf{x}, \tilde{\mathbf{x}} \in \mathcal{X}$. Mercer's theorem and various generalizations give conditions under which representations like (3)

are known to hold (Carmeli et al., 2006); one of the simplest examples is when \mathcal{X} is a compact Hausdorff space, $\rho_{\mathcal{X}}$ is a probability measure on the Borel sets of \mathcal{X} , and K is continuous. Observe that

$$\sum_{j=1}^{\infty} t_j^2 = \sum_{j=1}^{\infty} t_j^2 \int_{\mathcal{X}} \psi_j(\mathbf{x})^2 \, d\rho_{\mathcal{X}}(\mathbf{x}) = \int_{\mathcal{X}} K(\mathbf{x}, \mathbf{x}) \, d\rho_{\mathcal{X}}(\mathbf{x}) \leq \kappa^2;$$

in particular, $\{t_j^2\} \in \ell^1(\mathbb{N})$. Additionally, t_j^2 and ψ_j are eigenvalues and eigenfunctions, respectively, of the linear transformation $T : L^2(\rho_{\mathcal{X}}) \rightarrow L^2(\rho_{\mathcal{X}})$ defined by

$$(Tf)(\tilde{\mathbf{x}}) = \int_{\mathcal{X}} f(\mathbf{x})K(\mathbf{x}, \tilde{\mathbf{x}}) \, d\rho_{\mathcal{X}}(\mathbf{x}) = \sum_{j=1}^{\infty} t_j^2 \langle f, \psi_j \rangle_{L^2(\rho_{\mathcal{X}})} \psi_j(\tilde{\mathbf{x}}), \quad (4)$$

for $f \in L^2(\rho_{\mathcal{X}})$, $\tilde{\mathbf{x}} \in \mathcal{X}$, where $\langle \cdot, \cdot \rangle_{L^2(\rho_{\mathcal{X}})}$ denotes the inner product on $L^2(\rho_{\mathcal{X}})$.

Let $\mathcal{H} \subseteq L^2(\rho_{\mathcal{X}})$ be the RKHS corresponding to K (Aronszajn, 1950) and let $\phi_j = t_j \psi_j$, $j = 1, 2, \dots$. It follows from basic facts about RKHSs that $\{\phi_j\}_{j=1}^{\infty}$ is an orthonormal basis for \mathcal{H} (if $t_j^2 > t_{j+1}^2 = 0$, then $\{\phi_j\}_{j=1}^J$ is an orthonormal basis for \mathcal{H}). Furthermore, \mathcal{H} is characterized by

$$\mathcal{H} = \left\{ f = \sum_{j=1}^{\infty} \theta_j \psi_j \in L^2(\rho_{\mathcal{X}}); \sum_{j=1}^{\infty} \frac{\theta_j^2}{t_j^2} < \infty \right\}$$

and the inner product

$$\langle f, \tilde{f} \rangle_{\mathcal{H}} = \left\langle \sum_{j=1}^{\infty} \theta_j \psi_j, \sum_{j=1}^{\infty} \tilde{\theta}_j \psi_j \right\rangle_{\mathcal{H}} = \sum_{j=1}^{\infty} \frac{\theta_j \tilde{\theta}_j}{t_j^2}$$

(the corresponding norm is denoted by $\|\cdot\|_{\mathcal{H}}$). Our main assumption on the relationship between y , \mathbf{x} , and the kernel K is that

$$f^\dagger \in \mathcal{H}. \quad (5)$$

This is a regularity or smoothness assumption on f^\dagger . Many of the results in this paper can be modified, so that they apply to settings where $f^\dagger \notin \mathcal{H}$, by replacing f^\dagger with an appropriate projection of f^\dagger and including an approximation error term in the corresponding bounds. This approach leads to the study of *oracle inequalities* (Zhang, 2005; Zhang et al., 2013; Koltchinskii, 2006; Steinwart et al., 2009; Hsu et al., 2014), which we do not pursue in detail here. However, investigating oracle inequalities for general regularization operators may be of interest for future research, as most existing work focuses on ridge regularization.

Another interpretation of condition (5) is that it is a minimal regularity condition on f^\dagger for ensuring that f^\dagger can be estimated consistently using the kernel methods considered below. One key aspect of the upper bounds in Section 5 is

that they show f^\dagger can be estimated more efficiently, if it satisfies stronger regularity conditions (and if the regularization method used is sufficiently adaptable). A convenient way to formulate a collection of regularity conditions with varying strengths, which will be useful in the sequel, is as follows. For $\zeta \geq 0$, define the Hilbert space

$$\mathcal{H}_\zeta = \left\{ f = \sum_{j=1}^{\infty} \theta_j \psi_j \in L^2(\rho_{\mathcal{X}}); \sum_{j=1}^{\infty} \frac{\theta_j^2}{t_j^{2(1+\zeta)}} < \infty \right\}. \quad (6)$$

Then $\mathcal{H} = \mathcal{H}_0$ and $\mathcal{H}_{\zeta_2} \subseteq \mathcal{H}_{\zeta_1}$ whenever $\zeta_2 \geq \zeta_1 \geq 0$. The norm on \mathcal{H}_ζ is defined by $\|f\|_{\mathcal{H}_\zeta}^2 = \sum_{j=1}^{\infty} \frac{\theta_j^2}{t_j^{2(1+\zeta)}}$. Additionally, positive integers J define the finite-rank subspace

$$\mathcal{H}_J^\circ = \left\{ f = \sum_{j=1}^J \theta_j \psi_j \in L^2(\rho_{\mathcal{X}}); \theta_1, \dots, \theta_J \in \mathbb{R} \right\}. \quad (7)$$

We have the inclusion $\mathcal{H}_J^\circ \subseteq \mathcal{H}_{J+1}^\circ$ and, if $t_1^2, \dots, t_J^2 > 0$, then $\mathcal{H}_J^\circ \subseteq \mathcal{H}_\zeta$ for any $\zeta \geq 0$. In particular, it is clear from (6)–(7) that $f^\dagger \in \mathcal{H}_\zeta$ is a stronger regularity condition than $f^\dagger \in \mathcal{H}$ and that $f^\dagger \in \mathcal{H}_J^\circ$ is an even stronger condition (provided the t_j^2 are strictly positive). Conditions such as $f^\dagger \in \mathcal{H}_\zeta$ and $f^\dagger \in \mathcal{H}_J^\circ$ are known as *source conditions* elsewhere in the literature (e.g., Bauer et al., 2007; Caponnetto and Yao, 2010).

4. Regularization

As discussed in Section 1, our goal is find estimators \hat{f} that minimize the risk (1). In this paper, we focus on regularization-based estimators for f^\dagger . In order to precisely describe these estimators, we require some additional notation for various operators that will be of interest, and some basic definitions from regularization theory.

4.1. Finite-rank operators of interest

For $\mathbf{x} \in \mathcal{X}$, define $K_{\mathbf{x}} \in \mathcal{H}$ by $K_{\mathbf{x}}(\tilde{\mathbf{x}}) = K(\mathbf{x}, \tilde{\mathbf{x}})$, $\tilde{\mathbf{x}} \in \mathcal{X}$. Let $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times d}$ and $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$. Additionally, define the finite-rank linear operators $S_X : \mathcal{H} \rightarrow \mathbb{R}^n$ and $T_X : \mathcal{H} \rightarrow \mathcal{H}$ (both depending on X) by

$$\begin{aligned} S_X f &= (\langle f, K_{\mathbf{x}_1} \rangle_{\mathcal{H}}, \dots, \langle f, K_{\mathbf{x}_n} \rangle_{\mathcal{H}})^\top = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^\top, \\ T_X f &= \frac{1}{n} \sum_{i=1}^n \langle f, K_{\mathbf{x}_i} \rangle_{\mathcal{H}} K_{\mathbf{x}_i} = \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}_i) K_{\mathbf{x}_i}, \end{aligned} \quad (8)$$

where $f \in \mathcal{H}$. Let $\langle \cdot, \cdot \rangle_{\mathbb{R}^n}$ denote the normalized inner-product on \mathbb{R}^n , defined by $\langle \mathbf{v}, \tilde{\mathbf{v}} \rangle_{\mathbb{R}^n} = n^{-1} \mathbf{v}^\top \tilde{\mathbf{v}}$ for $\mathbf{v} = (v_1, \dots, v_n)^\top$, $\tilde{\mathbf{v}} = (\tilde{v}_1, \dots, \tilde{v}_n)^\top \in \mathbb{R}^n$. Then

the adjoint of S_X with respect to $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and $\langle \cdot, \cdot \rangle_{\mathbb{R}^n}$, $S_X^* : \mathbb{R}^n \rightarrow \mathcal{H}$, is given by $S_X^* \mathbf{v} = \frac{1}{n} \sum_{i=1}^n v_i K_{\mathbf{x}_i}$. Additionally, we have $T_X = S_X^* S_X$. Finally, observe that $S_X S_X^* : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is given by the $n \times n$ matrix $S_X S_X^* = n^{-1} \mathbf{K}$, where $\mathbf{K} = (K(\mathbf{x}_i, \mathbf{x}_j))_{1 \leq i, j \leq n}$; \mathbf{K} is the *kernel matrix*, which is ubiquitous in kernel methods and enables finite computation.

4.2. Basic definitions

A family of functions $g_\lambda : [0, \infty) \rightarrow [0, \infty)$ indexed by $\lambda > 0$ is called a *regularization family* if it satisfies the following three conditions:²

R1 $\sup_{0 < t \leq \kappa^2} |tg_\lambda(t)| < 1$.

R2 $\sup_{0 < t \leq \kappa^2} |1 - tg_\lambda(t)| \leq 1$.

R3 $\sup_{0 < t \leq \kappa^2} |g_\lambda(t)| < \lambda^{-1}$.

The main idea behind a regularization family is that it “looks” similar to $t \mapsto 1/t$, but is better-behaved near $t = 0$, i.e., it is bounded by λ^{-1} . An important quantity that is related to the adaptability of a regularization family is the *qualification* of the regularization. The qualification of the regularization family $\{g_\lambda\}_{\lambda > 0}$ is defined to be the maximal $\xi \geq 0$ such that

$$\sup_{0 < t \leq \kappa^2} |1 - tg_\lambda(t)| t^\xi \leq \lambda^\xi.$$

If a regularization family has qualification ξ , we say that it “saturates at ξ .” Two regularization families that are the major motivation for the results in this paper are ridge (Tikhonov) regularization, where $g_\lambda(t) = \frac{1}{t+\lambda}$, and principal component (spectral cut-off) regularization, where

$$g_\lambda(t) = \frac{1}{t} \mathbb{1}\{t \geq \lambda\}. \quad (9)$$

Observe that ridge regularization has qualification 1 and principal component regularization has qualification ∞ . Another example of a regularization family with qualification ∞ is the Landweber iteration, which can be viewed as a special case of gradient descent (see, e.g., Rosasco et al., 2005; Bauer et al., 2007; Lo Gerfo et al., 2008).

4.3. Estimators

Given a regularization family $\{g_\lambda\}_{\lambda > 0}$, we define the g_λ -regularized estimator for f^\dagger ,

$$\hat{f}_\lambda = g_\lambda(T_X) S_X^* \mathbf{y}. \quad (10)$$

Here, g_λ is applied to the spectrum of the self-adjoint finite-rank operator T_X via the functional calculus (see, e.g., Reed and Simon, 1980, Ch. 7); this

²This definition follows Engl et al. (1996) and Bauer et al. (2007), but is slightly more restrictive.

is a generalization of standard matrix functions (Tropp, 2015, Ch. 2). Since $S_X^* g_\lambda(\mathbf{K}/n) = g_\lambda(T_X) S_X^*$, we have the finitely computable representation $\hat{f}_\lambda = S_X^* g_\lambda(\mathbf{K}/n) \mathbf{y} = \sum_{i=1}^n \hat{\gamma}_i K_{\mathbf{x}_i}$, where $(\hat{\gamma}_1, \dots, \hat{\gamma}_n)^\top = g_\lambda(\mathbf{K}/n) \mathbf{y}/n$; computing the $\hat{\gamma}_i$ involves an eigenvalue decomposition of the matrix \mathbf{K} . The dependence of \hat{f}_λ on the regularization family is implicit; our results hold for any regularization family except where explicitly stated otherwise. The estimators \hat{f}_λ are the main focus of this paper.

5. Main results

5.1. General bound on the risk

Theorem 1. *Let \hat{f}_λ be the estimator defined in (10) with regularization family $\{g_\lambda\}_{\lambda>0}$. Let $0 \leq \delta \leq 1$ and assume that there is some $\kappa_\delta^2 > 0$ such that*

$$\sup_{\mathbf{x} \in \mathcal{X}} \sum_{j=1}^{\infty} t_j^{2(1-\delta)} \psi_j(\mathbf{x})^2 \leq \kappa_\delta^2 < \infty \tag{11}$$

Assume that the source condition $f^\dagger \in \mathcal{H}_\zeta$ holds for some $\zeta \geq 0$, and that g_λ has qualification at least $\max\{(\zeta + 1)/2, 1\}$. Define the effective dimension $d_\lambda = \sum_{j=1}^{\infty} \frac{t_j^2}{t_j^2 + \lambda}$. Finally, assume that $(8/3 + 2\sqrt{5/3})\kappa_\delta^2/n \leq \lambda^{1-\delta} \leq \kappa_\delta^2$. The following risk bound holds:

$$\begin{aligned} \mathcal{R}_\rho(\hat{f}_\lambda) &\leq 2^{\zeta+3} \|f^\dagger\|_{\mathcal{H}_\zeta}^2 \lambda^{\zeta+1} + \frac{4d_\lambda \sigma^2}{n} \\ &\quad + 4d_\lambda \left(\|f^\dagger\|_{\mathcal{H}}^2 t_1^2 + \frac{\kappa^2 \sigma^2}{\lambda n} \right) \exp\left(-\frac{3\lambda^{1-\delta} n}{28\kappa_\delta^2}\right) \\ &\quad + \mathbf{1}\{\zeta > 1\} \cdot 16\zeta^2 (3/2)^{\zeta-1} \cdot \|f^\dagger\|_{\mathcal{H}_\zeta}^2 (t_1^2 + \lambda)^{\zeta-1} \kappa^4 \left(\frac{34}{n} + \frac{15}{n^2}\right). \end{aligned} \tag{12}$$

Theorem 1 is proved in the Appendices. The first two terms in the upper bound (12) are typically the dominant terms. In the upper bound (12), the interaction between the kernel K and the distribution $\rho_{\mathcal{X}}$ is reflected in the effective dimension d_λ (see, e.g., Zhang, 2005; Caponnetto and De Vito, 2007). The regularity of f^\dagger enters through norm of f^\dagger (both the \mathcal{H} - and \mathcal{H}_ζ -norms) and the exponent on λ .

The condition (11) in Theorem 1 is always satisfied by taking $\delta = 0$ and $\kappa_0^2 = \kappa^2$. Requiring (11) with $\delta > 0$ imposes additional conditions on the RKHS \mathcal{H} . For Corollary 1 below (which applies when the eigenvalues $\{t_j^2\}$ have polynomial-decay), we take $\delta = 0$ and $\kappa_0^2 = \kappa^2$. The stronger condition with $\delta > 0$ is required to obtain obtain minimax rates for kernels where the eigenvalues $\{t_j^2\}$ have exponential or Gaussian-type decay (see Corollaries 2-3).

Risk bounds on general regularization estimators similar to \hat{f}_λ were previously obtained by Bauer et al. (2007). However, their bounds (e.g., Theorem

10 of Bauer et al., 2007) are independent of the ambient RKHS \mathcal{H} , i.e., they do not depend on the eigenvalues $\{t_j^2\}$. Our bounds are tighter than those of Bauer et al. because we take advantage of the structure of \mathcal{H} . In contrast with our Theorem 1, the results of Bauer et al. do not give minimax bounds (not easily, at least), because minimax rates must depend on the t_j^2 .

5.2. Implications for kernels characterized by their eigenvalues' rate of decay

We now state consequences of Theorem 1 that give explicit rates for estimating f^\dagger via \hat{f}_λ , for any regularization family, under specific assumptions about the decay rate of the eigenvalues $\{t_j^2\}$.

We first consider the case where the eigenvalues have polynomial decay.

Corollary 1. *Assume that $C \geq 0$ and $\nu > 1/2$ are constants such that $0 < t_j^2 \leq Cj^{-2\nu}$ for all $j = 1, 2, \dots$. Assume the source condition $f^\dagger \in \mathcal{H}_\zeta$ for some $\zeta \geq 0$, and that g_λ has qualification at least $\max\{(\zeta + 1)/2, 1\}$. Finally, take $\lambda = C'n^{-\frac{2\nu}{2\nu(\zeta+1)+1}}$ for a suitable constant $C' > 0$ so that the conditions on λ from Theorem 1 are satisfied. Then*

$$\mathcal{R}_\rho(\hat{f}_\lambda) = O\left(\left\{\|f^\dagger\|_{\mathcal{H}_\zeta}^2 + \sigma^2\right\} \cdot n^{-\frac{2\nu(\zeta+1)}{2\nu(\zeta+1)+1}}\right),$$

where the constants implicit in the big- O may depend on κ^2 , C , C' , ν , and ζ , but nothing else.

Remark 1. *Observe that if g_λ has qualification at least $\max\{(\zeta + 1)/2, 1\}$ (and the other conditions of Corollary 1 are met), then \hat{f}_λ obtains the minimax rate for estimating functions over \mathcal{H}_ζ (Pinsker, 1980). Thus, if g_λ has higher qualification, then \hat{f}_λ can effectively adapt to a broader range of subspaces $\mathcal{H}_\zeta \subseteq \mathcal{H} = \mathcal{H}_0$. In particular, KPCR (with infinite qualification) can adapt to source conditions with arbitrary $\zeta \geq 0$; on the other hand, KRR satisfies the conditions of Corollary 1 only when $\zeta \leq 1$, because KRR has qualification 1.*

Remark 2. *As mentioned earlier, a very similar result for polynomial decay eigenvalues was independently and simultaneously obtained by Blanchard and Mücke (2016) for essentially the same class of regularization operators. Our Theorem 1, from which the corollary follows, applies to a broader class of kernels than results by Blanchard and Mücke.*

When the eigenvalues $\{t_j^2\}$ have exponential or Gaussian-type decay, the rates are nearly the same as in finite dimensions.

Corollary 2. *Assume that $C, \alpha \geq 0$ are constants such that $0 < t_j^2 \leq Ce^{-\alpha j}$ for all $j = 1, 2, \dots$. Assume that g_λ has qualification at least 1 and that (11) holds for any $0 < \delta \leq 1$. Finally, take $\lambda = C'n^{-1} \log(n)$ for a suitable constant $C' > 0$ so that the conditions on λ from Theorem 1 are satisfied. Then*

$$\mathcal{R}_\rho(\hat{f}_\lambda) = O\left(\left\{\|f^\dagger\|_{\mathcal{H}}^2 + \sigma^2\right\} \cdot \frac{\log(n)}{n}\right),$$

where the constants implicit in the big-O may depend on κ^2 , C , C' , α , δ , and κ_δ^2 , but nothing else.

Corollary 3. *Assume that $C, \alpha \geq 0$ are constants such that $0 < t_j^2 \leq Ce^{-\alpha j^2}$ for all $j = 1, 2, \dots$. Assume that g_λ has qualification at least 1 and that (11) holds for any $0 < \delta \leq 1$. Finally, take $\lambda = C'n^{-1}\sqrt{\log(n)}$ for a suitable constant $C' > 0$ so that the conditions on λ from Theorem 1 are satisfied. Then*

$$\mathcal{R}_\rho(\hat{f}_\lambda) = O\left(\left\{\|f^\dagger\|_{\mathcal{H}}^2 + \sigma^2\right\} \cdot \frac{\sqrt{\log(n)}}{n}\right),$$

where the constants implicit in the big-O may depend on κ^2 , C , C' , α , δ , and κ_δ^2 , but nothing else.

Remark 3. *In Corollaries 2–3, we get minimax estimation over $\mathcal{H} = \mathcal{H}_0$ (Pinsker, 1980). However, our bounds are not refined enough to pick up any potential improvements which may be had if a stronger source condition is satisfied (e.g., $f^\dagger \in \mathcal{H}_\zeta$ for $\zeta > 0$). This is typical in settings like this because the minimax rate is already quite fast, i.e., within a log-factor of the parametric rate n^{-1} .*

5.3. Parametric rates for finite-dimensional kernels and subspaces

If the kernel has finite rank (i.e., $t_j^2 = 0$ for j sufficiently large), then it follows directly from Theorem 1 that $\mathcal{R}_\rho(\hat{f}_\lambda) = O(\{\|f^\dagger\|_{\mathcal{H}}^2 + \sigma^2\}/n)$. If the kernel has infinite rank, but f^\dagger is contained in the finite-dimensional subspace \mathcal{H}_J° for some $J < \infty$, then Theorem 1 can still be applied, provided g_λ has high qualification. Indeed, if g_λ has infinite qualification and $f^\dagger \in \mathcal{H}_J^\circ$, then it follows that $f^\dagger \in \mathcal{H}_\zeta$ for all $\zeta \geq 0$ and Theorem 1 implies that for any $0 < \alpha \leq 1$, $\mathcal{R}_\rho(\hat{f}_\lambda) = O(\{\|f^\dagger\|_{\mathcal{H}}^2 + \sigma^2\}/n^{1-\alpha})$ for appropriately chosen λ . In fact, we can improve on this rate for KPCR. The next proposition implies that the risk of KPCR matches the parametric rate n^{-1} for $f^\dagger \in \mathcal{H}_J^\circ$; the proof requires a different argument, based on eigenvalue perturbation theory (found in the Appendices).

Proposition 1. *Let $\hat{f}_{\text{KPCR},\lambda}$ be the KPCR estimator, with principal component regularization (9), and assume that $f^\dagger \in \mathcal{H}_J^\circ$. Let $0 < r < 1$ be a constant and let $\lambda = (1-r)t_J^2$. If $rt_J^2 \geq \kappa^2/n^{1/2} + \kappa^2/(3n)$, then*

$$\begin{aligned} \mathcal{R}_\rho(\hat{f}_{\text{KPCR},\lambda}) \leq & \frac{1}{n} \cdot \left\{ \frac{34\kappa^6}{r^2 t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 + \frac{3\kappa^2}{(1-r)t_J^2} \sigma^2 \right\} \\ & + \kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \left\{ \frac{15\kappa^4}{n^2 r^2 t_J^4} + 4 \exp\left(-\frac{nr^2 t_J^4}{2\kappa^4 + 2\kappa^2 r t_J^2 / 3}\right) \right\}. \end{aligned}$$

Proposition 1 implies that KPCR may reach the parametric rate for estimating $f^\dagger \in \mathcal{H}_J^\circ$. On the other hand, it is known that KRR may perform dramatically worse than KPCR in these settings due to the saturation effect (see, e.g., Caponnetto and De Vito, 2007; Dicker, 2016; Dhillon et al., 2013).

6. Numerical experiments

6.1. Simulated data

This simulation study shows how KPCR is able to adapt to highly structured signals, while KRR requires more favorable structure from the ambient RKHS.

For this experiment, we take $\mathcal{X} = \{1, 2, \dots, 2^{13}\}$. The data distribution ρ on $\mathcal{Y} \times \mathcal{X}$ is specified as follows. The marginal distribution on \mathcal{X} is $\rho_{\mathcal{X}}(x) \propto x^{-1/2}$, the regression function f^\dagger is given by $f^\dagger(x) = \sum_{j=1}^5 \mathbb{1}\{x = j\}$, and $\rho(\cdot | x)$ is normal with mean $f^\dagger(x)$ and variance $1/4$.

To compute \hat{f}_λ , we use the discrete kernel $K(x, \tilde{x}) = \mathbb{1}\{x = \tilde{x}\}$. For each $x \in \mathcal{X}$, the estimate of $f^\dagger(x)$ is based on an empirical average; regularization via KRR and KPCR has the effect of shrinking the empirical estimate of $f^\dagger(x)$ towards zero.

Using an iid sample of size $n = 2^{13}$, we compute \hat{f}_λ (either KRR or KPCR) for λ in a discrete grid of 2^{10} values uniformly spaced between 10^{-5} and 0.02 , and then choose the value of λ for which \hat{f}_λ has smallest *validation mean-squared error* $n^{-1} \sum_{i=1}^n (y_i^v - \hat{f}_\lambda(x_i^v))^2$, computed using a separate iid sample $\{(x_i^v, y_i^v)\}_{i=1}^n$ of size $n = 2^{13}$.

Figure 1(a) shows the validation-MSE of each \hat{f}_λ ; the plot of $\rho_{\mathcal{X}}\text{-MSE} \|\hat{f}^\dagger -$

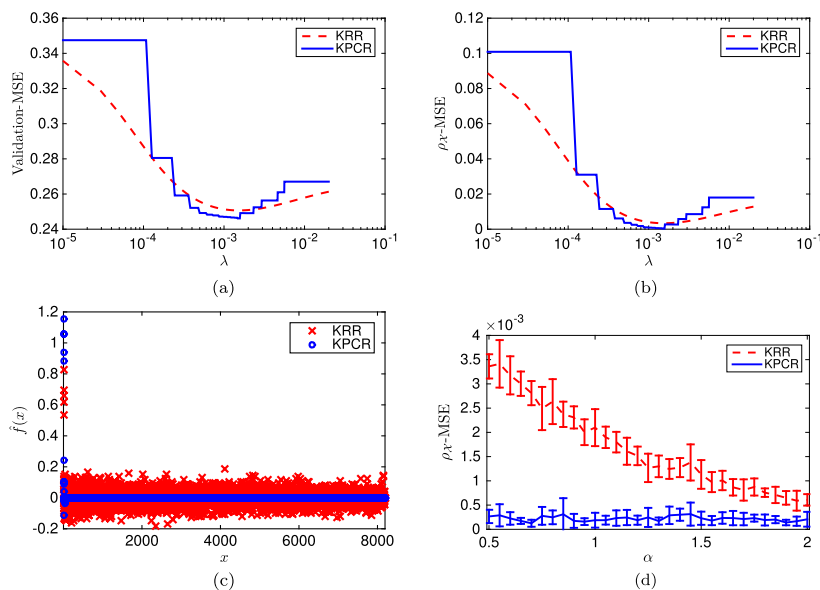


FIG 1. (a) Validation-MSE of $\hat{f}_{\text{KRR},\lambda}$ and $\hat{f}_{\text{KPCR},\lambda}$ for $\rho_{\mathcal{X}}(x) \propto x^{-1/2}$ vs. λ . (b) $\rho_{\mathcal{X}}\text{-MSE}$ of $\hat{f}_{\text{KRR},\lambda}$ and $\hat{f}_{\text{KPCR},\lambda}$ vs. λ with $\rho_{\mathcal{X}}(x) \propto x^{-1/2}$. (c) The estimated functions $\hat{f}_{\text{KRR},\hat{\lambda}_{\text{KRR}}}$ and $\hat{f}_{\text{KPCR},\hat{\lambda}_{\text{KPCR}}}$ from the same simulation. The regression function f^\dagger is $f^\dagger(x) = \sum_{j=1}^5 \mathbb{1}\{x = j\}$. The KRR function is non-zero for much of the domain, while the KPCR function is zero for nearly all of the domain (like f^\dagger). (d) $\rho_{\mathcal{X}}\text{-MSE}$ of $\hat{f}_{\text{KRR},\hat{\lambda}_{\text{KRR}}}$ and $\hat{f}_{\text{KPCR},\hat{\lambda}_{\text{KPCR}}}$ for $\rho_{\mathcal{X}}(x) \propto x^{-\alpha}$ vs. α .

$\hat{f}_\lambda \|_{L^2(\rho_X)}^2$ in Figure 1(b) has roughly the same shape, just shifted down by $\sigma^2 = 1/4$. The selected λ is $\hat{\lambda}_{\text{KRR}} = 0.001534$ for KRR, and $\hat{\lambda}_{\text{KPCR}} = 0.001573$ for KPCR. These choices of λ yield the final estimators, $\hat{f}_{\text{KRR}, \hat{\lambda}_{\text{KRR}}}$ and $\hat{f}_{\text{KPCR}, \hat{\lambda}_{\text{KPCR}}}$; the ρ_X -MSE is 0.0034 for KRR, and 0.0003 for KPCR. In Figure 1(c), we plot the functions $\hat{f}_{\text{KRR}, \hat{\lambda}_{\text{KRR}}}$ and $\hat{f}_{\text{KPCR}, \hat{\lambda}_{\text{KPCR}}}$ to show the effect of regularization. The KRR function is non-zero for much of the domain, while the KPCR function is zero for nearly all of the domain (like f^\dagger).

We repeat the above simulation for different marginal distributions $\rho_X(x) \propto x^{-\alpha}$, for $1/2 \leq \alpha \leq 2$, which imply different eigenvalue sequences $\{t_j^2\}$. The mean and standard deviation of the ρ_X -MSE's over 10 repetitions are shown in Figure 1(d). This confirms KPCR's to adapt to the regularity of f^\dagger , regardless of the ambient RKHS; KRR requires more structure to achieve similar results.

6.2. Real data

We also compared KRR and KPCR using three “weighted degree” kernels designed for recognizing splice sites in genetic sequences (Sonnenburg, 2008).³ The 3300 samples are divided into a training set (1000), validation set (1100), and testing set (1200). For each kernel, we use the training data to compute \hat{f}_λ for λ in a discrete grid of 2^{10} equally-spaced values between 10^{-5} and 0.4, and select the value of λ on which the MSE of \hat{f}_λ on the validation set is smallest. The MSE on the testing set and the intrinsic dimension $d_{\hat{\lambda}}$ for the selected $\hat{\lambda}$ (on the training data) are as follows:

	Kernel 1	Kernel 2	Kernel 3
MSE of $\hat{f}_{\text{KRR}, \hat{\lambda}_{\text{KRR}}}$	0.181452	0.172223	0.167745
MSE of $\hat{f}_{\text{KPCR}, \hat{\lambda}_{\text{KPCR}}}$	0.187059	0.168067	0.164159
$d_{\hat{\lambda}_{\text{KRR}}}$	175.6879	373.0029	738.0712
$d_{\hat{\lambda}_{\text{KPCR}}}$	170.9016	275.3560	581.1381

KRR outperforms KPCR with Kernel 1, where the intrinsic dimension of the kernel is low, while the reverse happens with Kernels 2 and 3, where the intrinsic dimensions are high. This resonates with our theoretical results, which suggest that KRR requires low intrinsic dimension to perform most effectively.

7. Discussion

Our unified analysis for a general class of regularization families in nonparametric regression highlights two important statistical properties. First, the results show minimax optimality for this general class in several commonly studied settings, which was only previously established for specific regularization methods. Second, the results demonstrate the adaptivity of certain regularization families

³We use the first three kernels for the data obtained from <http://mldata.org/repository/data/viewslug/mkl-splice/>.

to subspaces of the RKHS, showing that these techniques may take advantage of additional smoothness properties that the signal may possess. It is notable that the most well-studied family, KRR/Tikhonov regularization, does not possess this adaptability property.

Appendix A: Proof of Theorem 1

To provide some intuition behind \hat{f}_λ and our proof strategy, consider the operator T defined in (4), as applied to an element $f \in \mathcal{H}$,

$$Tf = \int_{\mathcal{X}} \langle f, K_{\mathbf{x}} \rangle_{\mathcal{H}} K_{\mathbf{x}} \, d\rho_{\mathcal{X}}(\mathbf{x}).$$

Observe that T is a “population” version of the operator T_X , defined in (8). Unlike T_X , T often has infinite rank; however, we still might expect that $T \approx T_X$ for large n , where the approximation holds in some suitable sense.

We also have a large- n approximation for $S_X^* \mathbf{y}$. For $f \in \mathcal{H}$,

$$\begin{aligned} \langle f, S_X^* \mathbf{y} \rangle_{\mathcal{H}} &= \frac{1}{n} \sum_{i=1}^n y_i \langle f, K_{\mathbf{x}_i} \rangle_{\mathcal{H}} \approx \int_{\mathcal{Y} \times \mathcal{X}} y f(\mathbf{x}) \, d\rho(y, \mathbf{x}) \\ &= \int_{\mathcal{X}} f^\dagger(\mathbf{x}) f(\mathbf{x}) \, d\rho_{\mathcal{X}}(\mathbf{x}) = \langle f, f^\dagger \rangle_{L^2(\rho_{\mathcal{X}})} = \langle f, T f^\dagger \rangle_{\mathcal{H}}, \end{aligned}$$

where we have used the identity $\phi_j = t_j \psi_j$ to obtain the last equality. It follows that $S_X^* \mathbf{y} \approx T f^\dagger$. Hence, to recover f^\dagger from \mathbf{y} , it would be natural to apply the inverse of T to $S_X^* \mathbf{y}$. However, T is not invertible whenever it has infinite rank, and regularization becomes necessary. We thus arrive at the chain of approximations which help motivate \hat{f}_λ :

$$\hat{f}_\lambda = g_\lambda(T_X) S_X^* \mathbf{y} \approx g_\lambda(T) T f^\dagger \approx f^\dagger,$$

where $g_\lambda(T)$ may be viewed as an approximate inverse for a suitably chosen regularization parameter λ .

A.1. Bias-variance decomposition

The proof of Theorem 1 is based on a simple bias-variance decomposition of the risk of \hat{f}_λ . Let $\boldsymbol{\epsilon} = (y_1 - f^\dagger(\mathbf{x}_1), \dots, y_n - f^\dagger(\mathbf{x}_n))^\top \in \mathbb{R}^n$.

Proposition 2. *The risk $\mathcal{R}_\rho(\hat{f}_\lambda)$ has the decomposition*

$$\mathcal{R}_\rho(\hat{f}_\lambda) = B_\rho(\hat{f}_\lambda) + V_\rho(\hat{f}_\lambda), \quad (13)$$

where $B_\rho(\hat{f}_\lambda)$ and $V_\rho(\hat{f}_\lambda)$ are defined as

$$\begin{aligned} B_\rho(\hat{f}_\lambda) &= \mathbb{E} \left[\left\| T^{1/2} \{I - g_\lambda(T_X) T_X\} f^\dagger \right\|_{\mathcal{H}}^2 \right], \\ V_\rho(\hat{f}_\lambda) &= \mathbb{E} \left[\left\| T^{1/2} g_\lambda(T_X) S_X^* \boldsymbol{\epsilon} \right\|_{\mathcal{H}}^2 \right]. \end{aligned}$$

Our proof separately bounds the bias $B_\rho(\hat{f}_\lambda)$ and variance $V_\rho(\hat{f}_\lambda)$ terms from Proposition 2. Taken together, these bounds imply a bound on $\mathcal{R}_\rho(\hat{f}_\lambda)$.

A.2. Translation to vector and matrix notation

We first note that the Hilbert space \mathcal{H} is isometric to $\ell^2(\mathbb{N})$ via the isometric isomorphism $\iota: \mathcal{H} \rightarrow \ell^2(\mathbb{N})$, given by

$$\iota: \sum_{j=1}^{\infty} \alpha_j \phi_j \mapsto (\alpha_1, \alpha_2, \dots)^\top \tag{14}$$

(we take all elements of $\ell^2(\mathbb{N})$ to be infinite-dimensional column vectors). Using this equivalence, we can convert elements of \mathcal{H} and linear operators on \mathcal{H} appearing in Proposition 2 into (infinite-dimensional) vectors and matrices, respectively, which we find simpler to analyze in the sequel.

Define the (infinite-dimensional) diagonal matrix

$$\mathbf{T} = \text{diag}(t_1^2, t_2^2, \dots),$$

and the vector

$$\boldsymbol{\beta} = (\beta_1, \beta_2, \dots)^\top \in \ell^2(\mathbb{N}),$$

where $\beta_i = \langle f^\dagger, \phi_i \rangle_{\mathcal{H}}$. Next, define the $n \times \infty$ (random) matrices

$$\begin{aligned} \boldsymbol{\Psi} &= (\psi_j(\mathbf{x}_i))_{1 \leq i \leq n; 1 \leq j < \infty}, \\ \boldsymbol{\Phi} &= \boldsymbol{\Psi} \mathbf{T} = (\phi_j(\mathbf{x}_i))_{1 \leq i \leq n; 1 \leq j < \infty}. \end{aligned}$$

Observe that

$$\begin{aligned} S_X &= \boldsymbol{\Phi} \circ \iota, \\ S_X^* &= \iota^{-1} \circ \left(\frac{1}{n} \boldsymbol{\Phi}^\top \right), \\ T &= \iota^{-1} \circ \mathbf{T} \circ \iota. \end{aligned}$$

Also, for $1 \leq i \leq n$, let $\phi_i \phi_i^\top$ be the $\infty \times \infty$ matrix whose (j, j') -th entry is $\phi_j(\mathbf{x}_i) \phi_{j'}(\mathbf{x}_i)$, and define

$$\boldsymbol{\Sigma} = \frac{1}{n} \sum_{i=1}^n \phi_i \phi_i^\top = \frac{1}{n} \boldsymbol{\Phi}^\top \boldsymbol{\Phi}.$$

Finally, let $\mathbf{I} = \text{diag}(1, 1, \dots)$, and let $\|\cdot\| = \|\cdot\|_{\ell^2(\mathbb{N})}$ denote the norm on $\ell^2(\mathbb{N})$. In these matrix and vector notations, the bias-variance decomposition from Proposition 2 translates to the following:

$$B_\rho(\hat{f}_\lambda) = \mathbb{E} \left[\left\| \mathbf{T}^{1/2} \{ \mathbf{I} - g_\lambda(\boldsymbol{\Sigma}) \boldsymbol{\Sigma} \} \boldsymbol{\beta} \right\|^2 \right],$$

$$V_\rho(\hat{f}_\lambda) = \frac{1}{n^2} \mathbb{E} \left[\left\| \mathbf{T}^{1/2} g_\lambda(\boldsymbol{\Sigma}) \boldsymbol{\Phi}^\top \boldsymbol{\epsilon} \right\|^2 \right].$$

The boundedness of the kernel implies

$$\begin{aligned} \text{tr}(\mathbf{T}) &\leq \kappa^2, \\ \|\boldsymbol{\phi}_i \boldsymbol{\phi}_i^\top\| &\leq \kappa^2, \\ \|\boldsymbol{\Sigma}\| &\leq \kappa^2, \end{aligned}$$

where the norms are the operator norms in $\ell^2(\mathbb{N})$.

A.3. Probabilistic bounds

For $0 < r < 1$, define the event

$$\mathcal{A}_r = \left\{ \left\| (\mathbf{T} + \lambda \mathbf{I})^{-1/2} (\boldsymbol{\Sigma} - \mathbf{T}) (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \right\| \geq r \right\}, \quad (15)$$

and let \mathcal{A}_r^c denote its complement. Our bounds on bias $B_\rho(\hat{f}_\lambda)$ and variance $V_\rho(\hat{f}_\lambda)$ are based on analysis in the event \mathcal{A}_r^c (for a constant r). Therefore, we also need to show that \mathcal{A}_r^c has large probability (equivalently, show that \mathcal{A}_r has small probability).

Lemma 1. *Assume that*

$$\sup_{\mathbf{x} \in \mathcal{X}} \sum_{j=1}^{\infty} t_j^{2(1-\delta)} \psi_j(\mathbf{x})^2 \leq \kappa_\delta^2 < \infty \quad (16)$$

for some $0 \leq \delta < 1$, $\kappa_\delta^2 > 0$. Further assume that $\lambda^{1-\delta} \leq \kappa_\delta^2$. If $r \geq \sqrt{\kappa_\delta^2 / (\lambda^{1-\delta} n)} + \kappa_\delta^2 / (3\lambda^{1-\delta} n)$, then

$$P(\mathcal{A}_r) \leq 4d_\lambda \exp\left(-\frac{\lambda^{1-\delta} n r^2}{2\kappa_\delta^2(1+r/3)}\right).$$

Proof. The proof is an application of Lemma 5. Define, for $1 \leq i \leq n$,

$$\mathbf{X}_i = \frac{1}{n} (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \{\boldsymbol{\phi}_i \boldsymbol{\phi}_i^\top - \mathbf{T}\} (\mathbf{T} + \lambda \mathbf{I})^{-1/2},$$

as well as $\mathbf{Y} = \sum_{i=1}^n \mathbf{X}_i$. We have

$$\mathbf{Y} = (\mathbf{T} + \lambda \mathbf{I})^{-1/2} (\boldsymbol{\Sigma} - \mathbf{T}) (\mathbf{T} + \lambda \mathbf{I})^{-1/2}.$$

It is clear that $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$. Observe that

$$\sum_{j=1}^{\infty} \frac{t_j^2}{t_j^2 + \lambda} \psi_j(\mathbf{x}_i)^2 \leq \max_{j \geq 1} \frac{t_j^{2\delta}}{t_j^2 + \lambda} \sum_{j=1}^{\infty} t_j^{2(1-\delta)} \psi_j(\mathbf{x}_i)^2$$

$$\leq \kappa_\delta^2 \max_{j \geq 1} \frac{t_j^{2\delta}}{t_j^2 + \lambda} \leq \frac{\kappa_\delta^2}{\lambda^{1-\delta}}, \tag{17}$$

where the last inequality uses the inequality of arithmetic and geometric means. Therefore, by the assumption $\lambda^{1-\delta} \leq \kappa_\delta^2$,

$$\begin{aligned} \|\mathbf{X}_i\| &\leq \frac{1}{n} \max \left\{ \|(\mathbf{T} + \lambda \mathbf{I})^{-1/2} \phi_i \phi_i^\top (\mathbf{T} + \lambda \mathbf{I})^{-1/2}\|, \right. \\ &\quad \left. \|(\mathbf{T} + \lambda \mathbf{I})^{-1/2} \mathbf{T} (\mathbf{T} + \lambda \mathbf{I})^{-1/2}\| \right\} \\ &= \frac{1}{n} \max \left\{ \sum_{j=1}^{\infty} \frac{t_j^2}{t_j^2 + \lambda} \psi_j(\mathbf{x}_i)^2, \max_{j \geq 1} \frac{t_j^2}{t_j^2 + \lambda} \right\} \leq \frac{\kappa_\delta^2}{\lambda^{1-\delta} n}. \end{aligned}$$

Moreover,

$$\begin{aligned} &\mathbb{E}[\mathbf{X}_i^2] \\ &= \frac{1}{n^2} \mathbb{E} \left[(\mathbf{T} + \lambda \mathbf{I})^{-1/2} \phi_i \phi_i^\top (\mathbf{T} + \lambda \mathbf{I})^{-1} \phi_i \phi_i^\top (\mathbf{T} + \lambda \mathbf{I})^{-1/2} - (\mathbf{T} + \lambda \mathbf{I})^{-2} \mathbf{T}^2 \right] \\ &= \frac{1}{n^2} \mathbb{E} \left[\left\{ \sum_{j=1}^{\infty} \frac{t_j^2 \psi_j(\mathbf{x}_i)^2}{t_j^2 + \lambda} \right\} (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \phi_i \phi_i^\top (\mathbf{T} + \lambda \mathbf{I})^{-1/2} - (\mathbf{T} + \lambda \mathbf{I})^{-2} \mathbf{T}^2 \right]. \end{aligned}$$

Combining this with (17) gives

$$\begin{aligned} \|\mathbb{E}[\mathbf{Y}^2]\| &\leq \frac{\kappa_\delta^2}{\lambda^{1-\delta} n} \left\| \mathbb{E} \left[(\mathbf{T} + \lambda \mathbf{I})^{-1/2} \phi_1 \phi_1^\top (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \right] \right\| \\ &= \frac{\kappa_\delta^2}{\lambda^{1-\delta} n} \|(\mathbf{T} + \lambda \mathbf{I})^{-1} \mathbf{T}\| \leq \frac{\kappa_\delta^2}{\lambda^{1-\delta} n} \end{aligned}$$

and

$$\begin{aligned} \text{tr}(\mathbb{E}[\mathbf{Y}^2]) &\leq \frac{\kappa_\delta^2}{\lambda^{1-\delta} n} \text{tr} \left(\mathbb{E} \left[(\mathbf{T} + \lambda \mathbf{I})^{-1/2} \phi_1 \phi_1^\top (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \right] \right) \\ &= \frac{\kappa_\delta^2}{\lambda^{1-\delta} n} \text{tr} \left((\mathbf{T} + \lambda \mathbf{I})^{-1} \mathbf{T} \right) = \frac{d_\lambda \kappa_\delta^2}{\lambda^{1-\delta} n}. \end{aligned}$$

Applying Lemma 5 with $V = R = \kappa_\delta^2/(\lambda^{1-\delta} n)$ and $D = d_\lambda$, gives

$$P(\|\mathbf{Y}\| \geq r) \leq 4d_\lambda \exp \left(-\frac{\lambda^{1-\delta} n r^2}{2\kappa_\delta^2(1+r/3)} \right). \quad \square$$

Lemma 2.

$$\mathbb{E}[\|\Sigma - \mathbf{T}\|^2] \leq \frac{34\kappa^4}{n} + \frac{15\kappa^4}{n^2}.$$

Proof. The proof is an application of Lemma 6. Define, for $1 \leq i \leq n$, $\mathbf{X}_i = \frac{1}{n}(\phi_i \phi_i^\top - \mathbf{T})$, and also define $\mathbf{Y} = \sum_{i=1}^n \mathbf{X}_i$, so $\mathbf{Y} = \Sigma - \mathbf{T}$. Clearly $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$ and $\|\mathbf{X}_i\| \leq \kappa^2/n$. Additionally, since $\mathbb{E}[\mathbf{Y}^2] = n^{-1} \mathbb{E}[\|\phi_1\|^2 \phi_1 \phi_1^\top - \mathbf{T}^2]$, we have $\|\mathbb{E}[\mathbf{Y}^2]\| \leq \kappa^4/n$ and $\text{tr}(\mathbb{E}[\mathbf{Y}^2]) \leq \kappa^4/n$. The claim thus follows by applying Lemma 6 with $V = \kappa^4/n$, $D = 1$, and $R = \kappa^2/n$. \square

A.4. Bias bound

Lemma 3. Assume that $\boldsymbol{\beta} = \mathbf{T}^{\zeta/2}\boldsymbol{\alpha}$ for some $\zeta \geq 0$ and that g_λ has qualification at least $\max\{(\zeta + 1)/2, 1\}$. For any $0 < r \leq 1/2$,

$$\begin{aligned} B_\rho(\hat{f}_\lambda) &\leq \frac{2^{\zeta+2}}{1-r} \|\boldsymbol{\alpha}\|^2 \lambda^{\zeta+1} + t_1^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_r) \\ &\quad + \mathbf{1}\{\zeta > 1\} \cdot \frac{8\zeta^2(1+r)^{\zeta-1}}{1-r} \cdot \|\boldsymbol{\alpha}\|^2 \cdot \|\mathbf{T} + \lambda\mathbf{I}\|^{\zeta-1} \cdot \mathbb{E}\left[\|\boldsymbol{\Sigma} - \mathbf{T}\|^2\right]. \end{aligned}$$

Proof. Define $h_\lambda(\boldsymbol{\Sigma}) = \mathbf{I} - \boldsymbol{\Sigma}g_\lambda(\boldsymbol{\Sigma})$. Since $|1 - sg_\lambda(s)| \leq 1$ for $0 < s \leq \kappa^2$ and $\|\boldsymbol{\Sigma}\| \leq \kappa^2$, we have

$$\|h_\lambda(\boldsymbol{\Sigma})\| \leq 1.$$

Moreover, using $\boldsymbol{\beta} = \mathbf{T}^{\zeta/2}\boldsymbol{\alpha}$,

$$\begin{aligned} B_\rho(\hat{f}_\lambda) &= \mathbb{E}\left[\left\|\mathbf{T}^{1/2}h_\lambda(\boldsymbol{\Sigma})\boldsymbol{\beta}\right\|^2\right] \\ &\leq \mathbb{E}\left[\left\|\mathbf{T}^{1/2}h_\lambda(\boldsymbol{\Sigma})\mathbf{T}^{\zeta/2}\right\|^2 \cdot \mathbf{1}\{\mathcal{A}_d^c\}\right] \cdot \|\boldsymbol{\alpha}\|^2 \\ &\quad + \|\mathbf{T}^{1/2}\|^2 \cdot \|h_\lambda(\boldsymbol{\Sigma})\|^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_d) \\ &\leq \mathbb{E}\left[\left\|\mathbf{T}^{1/2}h_\lambda(\boldsymbol{\Sigma})\mathbf{T}^{\zeta/2}\right\|^2 \cdot \mathbf{1}\{\mathcal{A}_d^c\}\right] \cdot \|\boldsymbol{\alpha}\|^2 + t_1^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_d). \end{aligned}$$

The rest of the proof involves bounding $\mathbb{E}[\|\mathbf{T}^{1/2}h_\lambda(\boldsymbol{\Sigma})\mathbf{T}^{\zeta/2}\|^2 \cdot \mathbf{1}\{\mathcal{A}_d^c\}]$. We separately consider two cases: (i) $\zeta \leq 1$, and (ii) $\zeta > 1$.

Case 1: $\zeta \leq 1$. Since g_λ has qualification at least 1, it follows that $|(s + \lambda)(1 - sg_\lambda(s))| \leq 2\lambda$ for $0 < s \leq \kappa^2$. This implies

$$\begin{aligned} \left\|\mathbf{T}^{1/2}h_\lambda(\boldsymbol{\Sigma})\mathbf{T}^{\zeta/2}\right\|^2 &\leq \left\|\mathbf{T}^{1/2}(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{-1/2}\right\|^2 \cdot \|(\boldsymbol{\Sigma} + \lambda\mathbf{I})h_\lambda(\boldsymbol{\Sigma})\|^2 \\ &\quad \cdot \left\|\mathbf{T}^{\zeta/2}(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{-1/2}\right\|^2 \\ &\leq 4\lambda^2 \cdot \left\|\mathbf{T}^{1/2}(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{-1}\mathbf{T}^{1/2}\right\| \\ &\quad \cdot \left\|\mathbf{T}^{\zeta/2}(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{-1}\mathbf{T}^{\zeta/2}\right\|. \end{aligned} \tag{18}$$

For $0 \leq z \leq 1$,

$$\begin{aligned} &\left\|\mathbf{T}^{z/2}(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{-1}\mathbf{T}^{z/2}\right\| \\ &= \left\|\mathbf{T}^{z/2}\{(\boldsymbol{\Sigma} - \mathbf{T}) + (\mathbf{T} + \lambda\mathbf{I})\}^{-1}\mathbf{T}^{z/2}\right\| \\ &\leq \left\|\mathbf{T}^{z/2}(\mathbf{T} + \lambda\mathbf{I})^{-1/2}\right\|^2 \cdot \left\|(\mathbf{T} + \lambda\mathbf{I})^{1/2}\{(\boldsymbol{\Sigma} - \mathbf{T}) + (\mathbf{T} + \lambda\mathbf{I})\}^{-1}(\mathbf{T} + \lambda\mathbf{I})^{1/2}\right\| \\ &= \left\|\mathbf{T}^z(\mathbf{T} + \lambda\mathbf{I})^{-1}\right\| \cdot \left\|\left(\mathbf{I} - (\mathbf{T} + \lambda\mathbf{I})^{-1/2}(\boldsymbol{\Sigma} - \mathbf{T})(\mathbf{T} + \lambda\mathbf{I})^{-1/2}\right)^{-1}\right\| \end{aligned}$$

$$\leq \lambda^{z-1} \cdot \left\| \left(\mathbf{I} - (\mathbf{T} + \lambda \mathbf{I})^{-1/2} (\mathbf{T} - \boldsymbol{\Sigma}) (\mathbf{T} + \lambda \mathbf{I})^{-1/2} \right)^{-1} \right\|, \quad (19)$$

where the final inequality uses the fact $s^z/(s+\lambda) \leq \lambda^{z-1}$ for $0 \leq z \leq 1$ and $s \geq 0$. The final quantity in (19) is bounded above by $\lambda^{z-1}/(1-r)$ on the event \mathcal{A}_r^c , so applying this inequality with $z = 1$ and $z = \zeta$ to (18) gives

$$\left\| \mathbf{T}^{1/2} h_\lambda(\boldsymbol{\Sigma}) \mathbf{T}^{\zeta/2} \right\|^2 \cdot \mathbf{1}_{\{\mathcal{A}_r^c\}} \leq 4\lambda^2 \cdot \frac{1}{1-r} \cdot \frac{\lambda^{\zeta-1}}{1-r} = \frac{4\lambda^{\zeta+1}}{(1-r)^2}.$$

So the bias in this case is bounded as

$$B_\rho(\hat{f}_\lambda) \leq \left(\frac{2}{1-r} \right)^2 \cdot \|\boldsymbol{\alpha}\|^2 \cdot \lambda^{\zeta+1} + t_1^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_r).$$

Case 2: $\zeta > 1$. We have

$$\begin{aligned} \left\| \mathbf{T}^{1/2} h_\lambda(\boldsymbol{\Sigma}) \mathbf{T}^{\zeta/2} \right\|^2 &\leq \left\| \mathbf{T}^{1/2} (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1/2} \right\|^2 \cdot \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma}) \mathbf{T}^{\zeta/2} \right\|^2 \\ &= \left\| \mathbf{T}^{1/2} (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1} \mathbf{T}^{1/2} \right\| \cdot \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma}) \mathbf{T}^{\zeta/2} \right\|^2 \\ &\leq \left\| \mathbf{T}^{1/2} (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1} \mathbf{T}^{1/2} \right\| \\ &\quad \cdot \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma}) (\mathbf{T} + \lambda \mathbf{I})^{\zeta/2} \right\|^2. \end{aligned} \quad (20)$$

The first factor on the right-hand side of (20) can be bounded using (19) on the event \mathcal{A}_r^c . For the second factor, we have

$$\begin{aligned} &\left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma}) (\mathbf{T} + \lambda \mathbf{I})^{\zeta/2} \right\| \\ &\leq \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{(\zeta+1)/2} h_\lambda(\boldsymbol{\Sigma}) \right\| \\ &\quad + \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma}) \left\{ (\mathbf{T} + \lambda \mathbf{I})^{\zeta/2} - (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{\zeta/2} \right\} \right\| \\ &\leq (2\lambda)^{(\zeta+1)/2} + 2\lambda^{1/2} \cdot \left\| (\mathbf{T} + \lambda \mathbf{I})^{\zeta/2} - (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{\zeta/2} \right\|. \end{aligned} \quad (21)$$

Above, the first inequality is due to the triangle inequality, and the second inequality uses the facts that g_λ has qualification at least $(\zeta+1)/2$, and that $(s+\lambda)^{(z+1)/2} \leq 2^{-1+(z+1)/2} (s^{(z+1)/2} + \lambda^{(z+1)/2})$ for $s \geq 0$ and $z \geq 1$.

We now bound $\|(\mathbf{T} + \lambda \mathbf{I})^{\zeta/2} - (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{\zeta/2}\|$ in terms of $\|\mathbf{T} - \boldsymbol{\Sigma}\|$. First, observe that on the event \mathcal{A}_r^c ,

$$\|\mathbf{T} - \boldsymbol{\Sigma}\| \leq \|\mathbf{T} + \lambda \mathbf{I}\| \cdot \|(\mathbf{T} + \lambda \mathbf{I})^{-1/2} (\mathbf{T} - \boldsymbol{\Sigma}) (\mathbf{T} + \lambda \mathbf{I})^{-1/2}\| < r \cdot \|\mathbf{T} + \lambda \mathbf{I}\|,$$

and, consequently,

$$\|\boldsymbol{\Sigma} + \lambda \mathbf{I}\| \leq (1+r) \cdot \|\mathbf{T} + \lambda \mathbf{I}\|.$$

For any constant $s > 0$, define

$$\begin{aligned} \mathbf{A}_s &= \frac{1}{(1+r+s)\|\mathbf{T} + \lambda\mathbf{I}\|} \cdot (\mathbf{T} + \lambda\mathbf{I}), \\ \mathbf{B}_s &= \frac{1}{(1+r+s)\|\mathbf{T} + \lambda\mathbf{I}\|} \cdot (\boldsymbol{\Sigma} + \lambda\mathbf{I}). \end{aligned}$$

Then, on the event \mathcal{A}_r^c , applying Lemma 7 and Lemma 8 gives

$$\begin{aligned} \|\mathbf{A}_s^{\zeta/2} - \mathbf{B}_s^{\zeta/2}\| &\leq 2\zeta \cdot \|\mathbf{A}_s^{1/2} - \mathbf{B}_s^{1/2}\| \\ &\leq \zeta \cdot \left\{ \frac{\lambda}{(1+r+s)\|\mathbf{T} + \lambda\mathbf{I}\|} \right\}^{-1/2} \cdot \|\mathbf{A}_s - \mathbf{B}_s\| \\ &= \zeta \cdot \{\lambda \cdot (1+r+s)\|\mathbf{T} + \lambda\mathbf{I}\|\}^{-1/2} \cdot \|\mathbf{T} - \boldsymbol{\Sigma}\|. \end{aligned}$$

Taking $s \rightarrow 0$, it follows that on \mathcal{A}_r^c ,

$$\begin{aligned} \|(\mathbf{T} + \lambda\mathbf{I})^{\zeta/2} - (\boldsymbol{\Sigma} + \lambda\mathbf{I})^{\zeta/2}\| &\leq \zeta \cdot \lambda^{-1/2} \cdot \{(1+r)\|\mathbf{T} + \lambda\mathbf{I}\|\}^{(\zeta-1)/2} \\ &\quad \cdot \|\mathbf{T} - \boldsymbol{\Sigma}\|. \end{aligned} \tag{22}$$

Combining (21) and (22) gives

$$\begin{aligned} \|(\boldsymbol{\Sigma} + \lambda\mathbf{I})^{1/2} h_\lambda(\boldsymbol{\Sigma})(\mathbf{T} + \lambda\mathbf{I})^{\zeta/2}\| \\ \leq (2\lambda)^{(\zeta+1)/2} + 2\zeta \cdot \{(1+r)\|\mathbf{T} + \lambda\mathbf{I}\|\}^{(\zeta-1)/2} \cdot \|\mathbf{T} - \boldsymbol{\Sigma}\|. \end{aligned}$$

Using this together with (19) in (20) and

$$\begin{aligned} &\left\| \mathbf{T}^{1/2} h_\lambda(\boldsymbol{\Sigma}) \mathbf{T}^{\zeta/2} \right\|^2 \cdot \mathbf{1}\{\mathcal{A}_r^c\} \\ &\leq \frac{1}{1-r} \cdot \left((2\lambda)^{(\zeta+1)/2} + 2\zeta \cdot \{(1+r)\|\mathbf{T} + \lambda\mathbf{I}\|\}^{(\zeta-1)/2} \cdot \|\mathbf{T} - \boldsymbol{\Sigma}\| \right)^2. \end{aligned}$$

Therefore

$$\begin{aligned} \mathbb{B}_\rho(\hat{f}_\lambda) &\leq \frac{\|\boldsymbol{\alpha}\|^2}{1-r} \cdot \left(2 \cdot (2\lambda)^{\zeta+1} + 8\zeta^2 \cdot \{(1+r)\|\mathbf{T} + \lambda\mathbf{I}\|\}^{\zeta-1} \cdot \mathbb{E}\left[\|\mathbf{T} - \boldsymbol{\Sigma}\|^2\right] \right) \\ &\quad + t_1^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_r). \end{aligned} \quad \square$$

A.5. Variance bound

Lemma 4. *For any $0 < r < 1$,*

$$\mathbb{V}_\rho(\hat{f}_\lambda) \leq \frac{2}{1-r} \cdot \frac{d_\lambda \sigma^2}{n} + \frac{\kappa^2 \sigma^2}{\lambda n} \cdot P(\mathcal{A}_r).$$

Proof. The assumption on $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)^\top$ implies $\mathbb{E}[\epsilon_i] = 0$, $\mathbb{E}[\epsilon_i \epsilon_j] = 0$ for $i \neq j$, and $\mathbb{E}[\epsilon_i^2] \leq \sigma^2$. So, by Von Neumann’s inequality,

$$\mathbb{V}_\rho(\hat{f}_\lambda) = \frac{1}{n^2} \mathbb{E} \left[\left\| \mathbf{T}^{1/2} g_\lambda(\boldsymbol{\Sigma}) \boldsymbol{\Phi}^\top \boldsymbol{\epsilon} \right\|^2 \right] \leq \frac{\sigma^2}{n} \mathbb{E} \left[\text{tr} \left(\mathbf{T} g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma} \right) \right].$$

Using Von Neumann's inequality together with $\text{tr}(\mathbf{T}) \leq \kappa^2$ and $g_\lambda(s)^2 s \leq 1/\lambda$ for $0 < s \leq \kappa^2$,

$$\text{tr}(\mathbf{T}g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma}) \leq \text{tr}(\mathbf{T}) \left\| g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma} \right\| \leq \frac{\kappa^2}{\lambda}.$$

Therefore

$$V_\rho(\hat{f}_\lambda) \leq \frac{\sigma^2}{n} \mathbb{E} \left[\text{tr}(\mathbf{T}g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma}) \cdot \mathbf{1}\{\mathcal{A}_r^c\} \right] + \frac{\kappa^2 \sigma^2}{\lambda n} \cdot P(\mathcal{A}_r).$$

Using Von Neumann's inequality twice more, and $(s + \lambda)g_\lambda(s)^2 s \leq 2$ for $0 < s \leq \kappa^2$, we have

$$\begin{aligned} \text{tr}(\mathbf{T}g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma}) &= \text{tr}(\mathbf{T}(\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1} (\boldsymbol{\Sigma} + \lambda \mathbf{I}) g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma}) \\ &\leq \text{tr}(\mathbf{T}(\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1}) \left\| (\boldsymbol{\Sigma} + \lambda \mathbf{I}) g_\lambda(\boldsymbol{\Sigma})^2 \boldsymbol{\Sigma} \right\| \\ &\leq 2 \text{tr}(\mathbf{T}(\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1}) \\ &= 2 \text{tr}(\mathbf{T}(\mathbf{T} + \lambda \mathbf{I})^{-1} (\mathbf{T} + \lambda \mathbf{I})^{1/2} (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1} (\mathbf{T} + \lambda \mathbf{I})^{1/2}) \\ &\leq 2 \text{tr}(\mathbf{T}(\mathbf{T} + \lambda \mathbf{I})^{-1}) \left\| (\mathbf{T} + \lambda \mathbf{I})^{1/2} (\boldsymbol{\Sigma} + \lambda \mathbf{I})^{-1} (\mathbf{T} + \lambda \mathbf{I})^{1/2} \right\| \\ &= 2d_\lambda \left\| \left(\mathbf{I} - (\mathbf{T} + \lambda \mathbf{I})^{1/2} (\mathbf{T} - \boldsymbol{\Sigma})^{-1} (\mathbf{T} + \lambda \mathbf{I})^{1/2} \right)^{-1} \right\|. \end{aligned}$$

This final quantity is bounded above by $2d_\lambda/(1-r)$ on the event \mathcal{A}_d^c . \square

A.6. Finishing the proof of Theorem 1

Using the bias-variance decomposition from Proposition 2, we apply the bias bound (Lemma 3) and variance bound (Lemma 4) to obtain a bound on the risk:

$$\begin{aligned} \mathcal{R}_\rho(\hat{f}_\lambda) &= B_\rho(\hat{f}_\lambda) + V_\rho(\hat{f}_\lambda) \\ &\leq \frac{2^{\zeta+2}}{1-r} \|\boldsymbol{\alpha}\|^2 \lambda^{\zeta+1} + t_1^2 \cdot \|\boldsymbol{\beta}\|^2 \cdot P(\mathcal{A}_r) \\ &\quad + \mathbf{1}\{\zeta > 1\} \cdot \frac{8\zeta^2(1+r)^{\zeta-1}}{1-r} \cdot \|\boldsymbol{\alpha}\|^2 \cdot (t_1^2 + \lambda)^{\zeta-1} \cdot \mathbb{E}[\|\boldsymbol{\Sigma} - \mathbf{T}\|^2] \\ &\quad + \frac{2}{1-r} \cdot \frac{d_\lambda \sigma^2}{n} + \frac{\kappa^2 \sigma^2}{\lambda n} \cdot P(\mathcal{A}_r). \end{aligned}$$

Now set $r = 1/2$, and apply Lemma 1 to bound $P(\mathcal{A}_r)$. Note that the assumption $(8/3 + 2\sqrt{5}/3)\kappa_\delta^2/n \leq \lambda^{1-\delta}$ satisfies the conditions on r in Lemma 1 for $r = 1/2$. Finally, apply Lemma 2 to bound $\mathbb{E}[\|\boldsymbol{\Sigma} - \mathbf{T}\|^2]$.

Appendix B: Proof of Proposition 1

Let $\hat{t}_1^2 \geq \hat{t}_2^2 \geq \dots \geq 0$ denote the eigenvalues of $\boldsymbol{\Sigma}$ and define $\hat{\mathbf{T}} = \text{diag}(\hat{t}_1^2, \hat{t}_2^2, \dots)$. Let \hat{U} be an orthogonal transformation satisfying $\boldsymbol{\Sigma} = \hat{U} \hat{\mathbf{T}} \hat{U}^\top$. Additionally, let

$h_\lambda(t) = I\{t \leq \lambda\}$, define $\hat{J} = \hat{J}_\lambda = \inf\{j; \hat{t}_j^2 > \lambda\}$, and write $\hat{U} = (\hat{U}_- \hat{U}_+)$, where \hat{U}_- is the $\infty \times \hat{J}$ matrix comprised of the first \hat{J} columns of \hat{U} , and \hat{U}_+ consists of the remaining columns of \hat{U} . Finally, define $\mathbf{T}_- = \text{diag}(t_1^2, \dots, t_{\hat{J}}^2)$, $\hat{\mathbf{T}}_+ = \text{diag}(\hat{t}_{\hat{J}+1}^2, \hat{t}_{\hat{J}+2}^2, \dots)$, and define the $\infty \times J$ matrix $U_- = (\mathbf{I}_J \ 0)^\top$, where \mathbf{I}_J is the $J \times J$ identity matrix.

Now consider the bias term $B_\rho(\hat{f}_{\text{KPCR},\lambda}) = \mathbb{E} \left[\left\| \mathbf{T}^{1/2} h_\lambda(\boldsymbol{\Sigma}) \boldsymbol{\beta} \right\|^2 \right]$, and observe that

$$\left\| \mathbf{T}^{1/2} h_\lambda(\boldsymbol{\Sigma}) \boldsymbol{\beta} \right\|^2 = \left\| \mathbf{T}^{1/2} \hat{U}_+ \hat{U}_+^\top U_- U_-^\top \boldsymbol{\beta} \right\|^2 \leq \kappa^2 \left\| \hat{U}_+^\top U_- \right\|^2 \|f^\dagger\|_{\mathcal{H}}^2.$$

Thus,

$$B_\rho(\hat{f}_{\text{KPCR},\lambda}) \leq \kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \mathbb{E} \left(\left\| \hat{U}_+^\top U_- \right\|^2 \cdot \mathbf{1}\{\hat{J} \geq J\} \right) + \kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 P(\hat{J} < J). \quad (23)$$

Now we bound $\|\hat{U}_+^\top U_- \|^2$ on the event $\{\hat{J} \geq J\}$. We derive this bound from basic principles, but it is essentially the Davis-Kahan inequality (Davis and Kahan, 1970). Let $\mathbf{D} = \boldsymbol{\Sigma} - \mathbf{T}$. Then

$$\mathbf{D}U_- = \boldsymbol{\Sigma}U_- - \mathbf{T}U_- = \boldsymbol{\Sigma}U_- - U_- \mathbf{T}_-$$

and

$$U_-^\top \mathbf{D} \hat{U}_+ = U_-^\top \boldsymbol{\Sigma} \hat{U}_+ - \mathbf{T}_- U_-^\top \hat{U}_+ = U_-^\top \hat{U}_+ \hat{\mathbf{T}}_+ - \mathbf{T}_- U_-^\top \hat{U}_+.$$

Next notice that

$$\begin{aligned} \|\mathbf{D}\| &\geq \|U_-^\top \mathbf{D} \hat{U}_+\| \geq \|\mathbf{T}_- U_-^\top \hat{U}_+\| - \|U_-^\top \hat{U}_+ \hat{\mathbf{T}}_+\| \\ &\geq t_J^2 \|U_-^\top \hat{U}_+\| - (1-r)t_J^2 \|U_-^\top \hat{U}_+\| = rt_J^2 \|U_-^\top \hat{U}_+\|. \end{aligned}$$

Thus,

$$\|U_-^\top \hat{U}_+\| \leq \frac{1}{rt_J^2} \|\mathbf{D}\| \quad (24)$$

on the event $\{\hat{J} > J\}$.

Next we bound $P(\hat{J} < J) = P(\hat{t}_J^2 \leq \lambda) = P\{\hat{t}_J^2 \leq (1-r)t_J^2\}$. By Weyl's inequality,

$$|\hat{t}_J^2 - t_J^2| \leq \|\mathbf{D}\| = \|\boldsymbol{\Sigma} - \mathbf{T}\|$$

and, by Lemma 5,

$$P\left(\|\boldsymbol{\Sigma} - \mathbf{T}\| \geq rt_J^2\right) \leq 4 \exp\left(-\frac{nr^2 t_J^4}{2\kappa^4 + 2\kappa^2 r t_J^2 / 3}\right),$$

provided $rt_J^2 \geq \kappa^2/n^{1/2} + \kappa^2/(3n)$. Thus,

$$P(\hat{J} < J) = P\left\{\hat{t}_J^2 \leq (1-r)t_J^2\right\}$$

$$\leq P(\|\mathbf{D}\| \geq rt_J^2) \leq 4 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right). \quad (25)$$

Combining (23)–(25) and using Lemma 2 gives

$$\begin{aligned} B_\rho(\hat{f}_{\text{KPCR},\lambda}) &\leq \frac{\kappa^2}{r^2t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 \mathbb{E} \left[\|\boldsymbol{\Sigma} - \mathbf{T}\|^2 \right] \\ &\quad + 4\kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right) \\ &\leq \frac{\kappa^2}{r^2t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 \left(\frac{34\kappa^4}{n} + \frac{15\kappa^4}{n^2} \right) \\ &\quad + 4\kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right). \end{aligned}$$

Next, we combine this bound on $B_\rho(\hat{f}_{\text{KPCR},\lambda})$ with the variance bound Lemma 4 (taking $r = 0$ in the lemma) to obtain

$$\begin{aligned} \mathcal{R}_\rho(\hat{f}_{\text{KPCR},\lambda}) &= B_\rho(\hat{f}_\lambda) + V_\rho(\hat{f}_\lambda) \\ &\leq \frac{\kappa^2}{r^2t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 \left(\frac{34\kappa^4}{n} + \frac{15\kappa^4}{n^2} \right) \\ &\quad + 4\kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right) + \frac{2d_\lambda\sigma^2}{n} + \frac{\kappa^2\sigma^2}{\lambda n} \\ &\leq \frac{\kappa^2}{r^2t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 \left(\frac{34\kappa^4}{n} + \frac{15\kappa^4}{n^2} \right) \\ &\quad + 4\kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right) + \frac{3\kappa^2\sigma^2}{(1-r)t_J^2n} \\ &= \frac{1}{n} \left\{ \frac{34\kappa^6}{r^2t_J^4} \|f^\dagger\|_{\mathcal{H}}^2 + \frac{3\kappa^2}{(1-r)t_J^2} \sigma^2 \right\} \\ &\quad + \kappa^2 \|f^\dagger\|_{\mathcal{H}}^2 \left\{ \frac{15\kappa^4}{n^2r^2t_J^4} + 4 \exp\left(-\frac{nr^2t_J^4}{2\kappa^4 + 2\kappa^2rt_J^2/3}\right) \right\}. \end{aligned}$$

This completes the proof of the proposition.

Appendix C: Additional proofs and tools

C.1. Sums of random operators

Lemma 5. *Let $R > 0$ be a positive real constant and consider a finite sequence of self-adjoint Hilbert-Schmidt operators $\{\mathbf{X}_i\}_{i=1}^n$ satisfying $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$*

and $\|\mathbf{X}_i\| \leq R$ almost surely. Define $\mathbf{Y} = \sum_{i=1}^n \mathbf{X}_i$ and suppose there are constants $V, D > 0$ satisfying $\|\mathbb{E}[\mathbf{Y}^2]\| \leq V$ and $\text{tr}(\mathbb{E}[\mathbf{Y}^2]) \leq VD$. For all $t \geq V^{1/2} + R/3$,

$$P(\|\mathbf{Y}\| \geq t) \leq 4D \exp\left(-\frac{t^2}{2(V + Rt/3)}\right).$$

Proof. This is a straightforward generalization of Tropp (2015, Theorem 7.7.1), using the arguments from Minsker (2011, Section 4) to extend from self-adjoint matrices to self-adjoint Hilbert-Schmidt operators. \square

Lemma 6. *In the same setting as Lemma 5,*

$$\mathbb{E}[\|\mathbf{Y}\|^2] \leq (2 + 32D)V + \left(\frac{2 + 128D}{9}\right)R^2.$$

Proof. The proof is based on integrating the tail bound from Lemma 5:

$$\begin{aligned} \mathbb{E}[\|\mathbf{Y}\|^2] &= \int_0^\infty P(\|\mathbf{Y}\| \geq t^{1/2}) dt \\ &\leq \left(V^{1/2} + \frac{R}{3}\right)^2 + \int_{(V^{1/2} + \frac{R}{3})^2}^\infty 4D \exp\left(-\frac{t}{2(V + Rt^{1/2}/3)}\right) dt \\ &\leq \left(V^{1/2} + \frac{R}{3}\right)^2 + \int_0^{9V^2/R^2} 4D \exp\left(-\frac{t}{4V}\right) dt \\ &\quad + \int_{9V^2/R^2}^\infty 4D \exp\left(-\frac{3t^{1/2}}{4R}\right) dt \\ &= \left(V^{1/2} + \frac{R}{3}\right)^2 + 16VD \left\{1 - \exp\left(-\frac{9V}{4R^2}\right)\right\} \\ &\quad + \frac{128R^2D}{9} \left\{\frac{9V}{4R^2} + 1\right\} \exp\left(-\frac{9V}{4R^2}\right) \\ &= \left(V^{1/2} + \frac{R}{3}\right)^2 + 16VD \left\{1 + \exp\left(-\frac{9V}{4R^2}\right)\right\} + \frac{128R^2D}{9} \exp\left(-\frac{9V}{4R^2}\right) \\ &\leq (2 + 32D)V + \left(\frac{2 + 128D}{9}\right)R^2. \square \end{aligned}$$

C.2. Differences of powers of bounded operators

Lemma 7. *Let \mathbf{A} and \mathbf{B} be non-negative self-adjoint operators with $\|\mathbf{A}\| < 1$ and $\|\mathbf{B}\| < 1$. For any $\gamma \geq 1$,*

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| \leq 2\gamma\|\mathbf{A} - \mathbf{B}\|.$$

Proof. The proof considers three possible cases for the value of γ : (i) γ is an integer, (ii) $1 < \gamma < 2$, and (iii) γ is a non-integer larger than two.

Case 1: γ is an integer. In this case, we have the following identity:

$$\mathbf{A}^\gamma - \mathbf{B}^\gamma = \sum_{j=1}^{\gamma} \mathbf{A}^{j-1}(\mathbf{A} - \mathbf{B})\mathbf{B}^{\gamma-j}.$$

So, by the triangle inequality,

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| \leq \|\mathbf{A} - \mathbf{B}\| \sum_{j=1}^{\gamma} \|\mathbf{A}^{j-1}\| \|\mathbf{B}^{\gamma-j}\| \leq \gamma \|\mathbf{A} - \mathbf{B}\|.$$

Case 2: $1 < \gamma < 2$. Pick $0 < r < 1$ such that $\|\mathbf{A}\| < r$ and $\|\mathbf{B}\| < r$, and fix $0 < t < (1 - r)/2$. Define $\mathbf{A}_t = \mathbf{A} + t\mathbf{I}$ and $\mathbf{B}_t = \mathbf{B} + t\mathbf{I}$, and

$$\binom{\gamma}{k} = \frac{\gamma(\gamma - 1) \cdots (\gamma - k + 1)}{k!}, \quad k = 1, 2, \dots.$$

Then, using the power series $(1 + s)^\gamma = 1 + \sum_{k=1}^{\infty} \binom{\gamma}{k} s^k$ for $-1 < s < 1$,

$$\mathbf{A}_t^\gamma - \mathbf{B}_t^\gamma = \sum_{k=1}^{\infty} \binom{\gamma}{k} \{(\mathbf{A}_t - \mathbf{I})^k - (\mathbf{B}_t - \mathbf{I})^k\} = \sum_{k=1}^{\infty} \binom{\gamma}{k} \sum_{j=1}^k \mathbf{A}^{j-1}(\mathbf{A} - \mathbf{B})\mathbf{B}^{k-j}.$$

Convergence is assured because $\|\mathbf{A}_t - \mathbf{I}\| \leq 1 - t$ and $\|\mathbf{B}_t - \mathbf{I}\| \leq 1 - t$. Moreover,

$$\begin{aligned} \|\mathbf{A}_t^\gamma - \mathbf{B}_t^\gamma\| &\leq \|\mathbf{A} - \mathbf{B}\| \sum_{k=1}^{\infty} \left| \binom{\gamma}{k} \right| \sum_{j=1}^k \|\mathbf{A}_t - \mathbf{I}\|^{j-1} \|\mathbf{B}_t - \mathbf{I}\|^{k-j} \\ &\leq \|\mathbf{A} - \mathbf{B}\| \sum_{k=1}^{\infty} \left| \binom{\gamma}{k} \right| k(1 - t)^{k-1} \\ &= \gamma \|\mathbf{A} - \mathbf{B}\| \left\{ 1 + \sum_{k=1}^{\infty} \left| \binom{\gamma - 1}{k} \right| (1 - t)^k \right\} \\ &= \gamma \|\mathbf{A} - \mathbf{B}\| \left\{ 2 - (1 - (1 - t))^{\gamma-1} \right\} \leq 2\gamma \|\mathbf{A} - \mathbf{B}\|. \end{aligned}$$

Above, we have used the power series $1 - (1 - s)^{\gamma-1} = \sum_{k=1}^{\infty} \left| \binom{\gamma-1}{k} \right| s^k$ at $s = 1 - t$. Taking $t \rightarrow 0$ yields

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| \leq 2\gamma \|\mathbf{A} - \mathbf{B}\|.$$

Case 3: γ is a non-integer larger than two. We can write $\gamma = k + q$ for an integer $k \geq 2$ and real number $0 < q < 1$. Applying the results from the previous two cases gives

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| = \|(\mathbf{A}^k)^{\gamma/k} - (\mathbf{B}^k)^{\gamma/k}\| \leq \frac{2\gamma}{k} \|\mathbf{A}^k - \mathbf{B}^k\| \leq 2\gamma \|\mathbf{A} - \mathbf{B}\|. \quad \square$$

Lemma 8. *Pick any real numbers $r, \gamma \in (0, 1)$. Let \mathbf{A} and \mathbf{B} be non-negative self-adjoint operators, each with spectrum contained in $[r, 1)$. Then*

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| \leq \gamma r^{\gamma-1} \|\mathbf{A} - \mathbf{B}\|.$$

Proof. Since $\|\mathbf{A} - \mathbf{I}\| \leq 1 - r$ and $\|\mathbf{B} - \mathbf{I}\| \leq 1 - r$, the proof is similar to that of Lemma 7:

$$\|\mathbf{A}^\gamma - \mathbf{B}^\gamma\| \leq \gamma \|\mathbf{A} - \mathbf{B}\| \left\{ 1 + \sum_{k=1}^{\infty} \left| \binom{\gamma-1}{k} \right| (1-r)^k \right\} = \gamma r^{\gamma-1} \|\mathbf{A} - \mathbf{B}\|.$$

Above, we have used the power series $(1-s)^{\gamma-1} = 1 + \sum_{k=1}^{\infty} \left| \binom{\gamma-1}{k} \right| s^k$ at $s = 1-r$ (which differs from the case in Lemma 7 because $0 < \gamma < 1$). \square

References

- ARONSAJN, N. (1950). Theory of reproducing kernels. *T. A. Math. Soc.* **68** 337–404.
- BAUER, F., PEREVERZEV, S. and ROSASCO, L. (2007). On regularization algorithms in learning theory. *J. Complexity* **23** 52–72.
- BLANCHARD, G. and MÜCKE, N. (2016). Optimal rates for regularization of statistical inverse learning problems. *arXiv preprint arXiv:1604.04054*.
- CAPONNETTO, A. and DE VITO, E. (2007). Optimal rates for the regularized least-squares algorithm. *Found. Comput. Math.* **7** 331–368. [MR2335249](#)
- CAPONNETTO, A. and YAO, Y. (2010). Cross-validation based adaptation for regularization operators in learning theory. *Anal. Appl.* **8** 161–183. [MR2646909](#)
- CARMELI, C., DE VITO, E. and TOIGO, A. (2006). Vector valued reproducing kernel Hilbert spaces of integrable functions and Mercer theorem. *Anal. Appl.* **4** 377–408.
- DAVIS, C. and KAHAN, W. M. (1970). The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* **7** 1–46. [MR0264450](#)
- DHILLON, P. S., FOSTER, D. P., KAKADE, S. M. and UNGAR, L. H. (2013). A risk comparison of ordinary least squares vs ridge regression. *J. Mach. Learn. Res.* **14** 1505–1511.
- DICKER, L. H. (2016). Ridge regression and asymptotic minimax estimation over spheres of growing dimension. *Bernoulli* **22** 1–37.
- ENGL, H. W., HANKE, M. and NEUBAUER, A. (1996). *Regularization of Inverse Problems*. Mathematics and Its Applications, Vol. 375, Springer.
- GUO, Z. C., LIN, S. B. and ZHOU, D. X. (2016). Learning theory of distributed spectral algorithms. Preprint.
- GYÖRFI, L., KOHLER, M., KRZYŻAK, A. and WALK, H. (2002). *A Distribution Free Theory of Nonparametric Regression*. Springer.
- HSU, D., KAKADE, S. M. and ZHANG, T. (2014). Random design analysis of ridge regression. *Found. Comput. Math.* **14** 569–600. [MR3201956](#)

- KOLTCHINSKII, V. (2006). Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Stat.* **34** 2593–2656. [MR2329442](#)
- LIN, S. B., GUO, X. and ZHOU, D. X. (2016). Distributed learning with regularized least squares. *arXiv preprint arXiv:1608.03339*.
- LO GERFO, L., ROSASCO, L., ODOE, F., DE VITO, E. and VERRI, A. (2008). Spectral algorithms for supervised learning. *Neural Comput.* **7** 1873–1897.
- MATHÉ, P. (2005). Saturation of regularization methods for linear ill-posed problems in Hilbert spaces. *SIAM J. Numer. Anal.* **42** 968–973.
- MINSKER, S. (2011). On some extensions of Bernstein’s inequality for self-adjoint operators. *arXiv preprint arXiv:1112.5448*.
- NEUBAUER, A. (1997). On converse and saturation results for Tikhonov regularization of linear ill-posed problems. *SIAM J. Numer. Anal.* **34** 517–527.
- PINSKER, M. S. (1980). Optimal filtration of square-integrable signals in Gaussian noise. *Probl. Inf. Transm.* **16** 52–68.
- REED, M. and SIMON, B. (1980). *Methods of Modern Mathematical Physics; Volume 1, Functional Analysis*. Academic Press.
- ROSASCO, L., DE VITO, E. and VERRI, A. (2005). Spectral methods for regularization in learning theory. Tech. Rep. DISI-TR-05-18, Università degli Studi di Genova, Italy.
- SONNENBURG, S. (2008). *Machine Learning for Genomic Sequence Analysis*. Ph.D. thesis, Fraunhofer Institute FIRST. Supervised by K.-R. Müller and G. Rätsch.
- STEINWART, I., HUSH, D. and SCOVEL, C. (2009). Optimal rates for regularized least squares regression. In *Conference on Learning Theory*.
- TROPP, J. A. (2015). An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*.
- WASSERMAN, L. (2006). *All of Nonparametric Statistics*. Springer.
- ZHANG, T. (2005). Learning bounds for kernel regression using effective data dimensionality. *Neural Comput.* **17** 2077–2098.
- ZHANG, Y., DUCHI, J. C. and WAINWRIGHT, M. J. (2013). Divide and conquer kernel ridge regression. In *Conference on Learning Theory*.