

Minimax-Optimal Rates For Sparse Additive Models Over Kernel Classes Via Convex Programming

Garvesh Raskutti
Martin J. Wainwright*
Bin Yu*

Department of Statistics
University of California
Berkeley, CA 94720-1776, USA

GARVESH@STAT.BERKELEY.EDU
 WAINWRIG@STAT.BERKELEY.EDU
 BINYU@STAT.BERKELEY.EDU

Editor: Nicolas Vayatis

Abstract

Sparse additive models are families of d -variate functions with the additive decomposition $f^* = \sum_{j \in S} f_j^*$, where S is an unknown subset of cardinality $s \ll d$. In this paper, we consider the case where each univariate component function f_j^* lies in a reproducing kernel Hilbert space (RKHS), and analyze a method for estimating the unknown function f^* based on kernels combined with ℓ_1 -type convex regularization. Working within a high-dimensional framework that allows both the dimension d and sparsity s to increase with n , we derive convergence rates in the $L^2(\mathbb{P})$ and $L^2(\mathbb{P}_n)$ norms over the class $\mathcal{F}_{d,s,\mathcal{H}}$ of sparse additive models with each univariate function f_j^* in the unit ball of a univariate RKHS with bounded kernel function. We complement our upper bounds by deriving minimax lower bounds on the $L^2(\mathbb{P})$ error, thereby showing the optimality of our method. Thus, we obtain optimal minimax rates for many interesting classes of sparse additive models, including polynomials, splines, and Sobolev classes. We also show that if, in contrast to our univariate conditions, the d -variate function class is assumed to be globally bounded, then much faster estimation rates are possible for any sparsity $s = \Omega(\sqrt{n})$, showing that global boundedness is a significant restriction in the high-dimensional setting.

Keywords: sparsity, kernel, non-parametric, convex, minimax

1. Introduction

The past decade has witnessed a flurry of research on sparsity constraints in statistical models. Sparsity is an attractive assumption for both practical and theoretical reasons: it leads to more interpretable models, reduces computational cost, and allows for model identifiability even under high-dimensional scaling, where the dimension d exceeds the sample size n . While a large body of work has focused on sparse linear models, many applications call for the additional flexibility provided by non-parametric models. In the general setting, a non-parametric regression model takes the form $y = f^*(x_1, \dots, x_d) + w$, where $f^* : \mathbb{R}^d \rightarrow \mathbb{R}$ is the unknown regression function, and w is scalar observation noise. Unfortunately, this general non-parametric model is known to suffer severely from the so-called ‘‘curse of dimensionality’’, in that for most natural function classes (e.g., twice differentiable functions), the sample size n required to achieve any given error grows exponentially in the dimension d . Given this curse of dimensionality, it is essential to further constrain the com-

*. Also in the department of Electrical Engineering & Computer Science.

plexity of possible functions f^* . One attractive candidate is the class of *additive non-parametric models* (see Hastie and Tibshirani, 1986), in which the function f^* has an additive decomposition of the form

$$f^*(x_1, x_2, \dots, x_d) = \sum_{j=1}^d f_j^*(x_j), \quad (1)$$

where each component function f_j^* is univariate. Given this additive form, this function class no longer suffers from the exponential explosion in sample size of the general non-parametric model. Nonetheless, one still requires a sample size $n \gg d$ for consistent estimation; note that this is true even for the linear model, which is a special case of Equation (1).

A natural extension of sparse linear models is the class of *sparse additive models*, in which the unknown regression function is assumed to have a decomposition of the form

$$f^*(x_1, x_2, \dots, x_d) = \sum_{j \in S} f_j^*(x_j), \quad (2)$$

where $S \subseteq \{1, 2, \dots, d\}$ is some unknown subset of cardinality $|S| = s$. Of primary interest is the case when the decomposition is genuinely sparse, so that $s \ll d$. To the best of our knowledge, this model class was first introduced by Lin and Zhang (2006), and has since been studied by various researchers (e.g., Koltchinskii and Yuan, 2010; Meier et al., 2009; Ravikumar et al., 2009; Yuan, 2007). Note that the sparse additive model (2) is a natural generalization of the sparse linear model, to which it reduces when each univariate function is constrained to be linear.

In past work, several groups have proposed computationally efficient methods for estimating sparse additive models (2). Just as ℓ_1 -based relaxations such as the Lasso have desirable properties for sparse parametric models, more general ℓ_1 -based approaches have proven to be successful in this setting. Lin and Zhang (2006) proposed the COSSO method, which extends the Lasso to cases where the component functions f_j^* lie in a reproducing kernel Hilbert space (RKHS); see also Yuan (2007) for a similar extension of the non-negative garrote (Breiman, 1995). Bach (2008) analyzes a closely related method for the RKHS setting, in which least-squares loss is penalized by an ℓ_1 -sum of Hilbert norms, and establishes consistency results in the classical (fixed d) setting. Other related ℓ_1 -based methods have been proposed in independent work by Koltchinskii and Yuan (2008), Ravikumar et al. (2009) and Meier et al. (2009), and analyzed under high-dimensional scaling ($d \gg n$). As we describe in more detail in Section 3.4, each of the above papers establish consistency and convergence rates for the prediction error under certain conditions on the covariates as well as the sparsity s and dimension d . However, it is not clear whether the rates obtained in these papers are sharp for the given methods, nor whether the rates are minimax-optimal. Past work by Koltchinskii and Yuan (2010) establishes rates for sparse additive models with an additional global boundedness condition, but as will be discussed at more length in the sequel, these rates are not minimax optimal in general.

This paper makes three main contributions to this line of research. Our first contribution is to analyze a simple polynomial-time method for estimating sparse additive models and provide upper bounds on the error in the $L^2(\mathbb{P})$ and $L^2(\mathbb{P}_n)$ norms. The estimator¹ that we analyze is based on a combination of least-squares loss with two ℓ_1 -based sparsity penalty terms, one corresponding to

1. The same estimator was proposed concurrently by Koltchinskii and Yuan (2010); we discuss connections to this work in the sequel.

an $\ell_1/L^2(\mathbb{P}_n)$ norm and the other an $\ell_1/\|\cdot\|_{\mathcal{H}}$ norm. Our first main result (Theorem 1) shows that with high probability, if we assume the univariate functions are bounded and independent, the error of our procedure in the squared $L^2(\mathbb{P}_n)$ and $L^2(\mathbb{P})$ norms is bounded by $O\left(\frac{s \log d}{n} + s v_n^2\right)$, where the quantity v_n^2 corresponds to the optimal rate for estimating a single univariate function. Importantly, our analysis does *not* require a global boundedness condition on the class $\mathcal{F}_{d,s,\mathcal{H}}$ of all s -sparse models, an assumption that is often imposed in classical non-parametric analysis. Indeed, as we discuss below, when such a condition is imposed, then significantly faster rates of estimation are possible. The proof of Theorem 1 involves a combination of techniques for analyzing M -estimators with decomposable regularizers (Negahban et al., 2009), combined with various techniques in empirical process theory for analyzing kernel classes (e.g., Bartlett et al., 2005; Mendelson, 2002; van de Geer, 2000). Our second contribution is complementary in nature, in that it establishes algorithm-independent minimax lower bounds on $L^2(\mathbb{P})$ error. These minimax lower bounds, stated in Theorem 2, are specified in terms of the metric entropy of the underlying univariate function classes. For both finite-rank kernel classes and Sobolev-type classes, these lower bounds match our achievable results, as stated in Corollaries 1 and 2, up to constant factors in the regime of sub-linear sparsity ($s = o(d)$). Thus, for these function classes, we have a sharp characterization of the associated minimax rates. The lower bounds derived in this paper initially appeared in the Proceedings of the NIPS Conference (December 2009). The proofs of Theorem 2 is based on characterizing the packing entropies of the class of sparse additive models, combined with classical information theoretic techniques involving Fano’s inequality and variants (e.g., Has’minskii, 1978; Yang and Barron, 1999; Yu, 1996).

Our third contribution is to determine upper bounds on minimax $L^2(\mathbb{P})$ and $L^2(\mathbb{P}_n)$ error when we impose a global boundedness assumption on the class $\mathcal{F}_{d,s,\mathcal{H}}$. More precisely, a global boundedness condition means that the quantity $B(\mathcal{F}_{d,s,\mathcal{H}}) = \sup_{f \in \mathcal{F}_{d,s,\mathcal{H}}} \sup_x |\sum_{j=1}^d f_j(x_j)|$ is assumed to be bounded independently of (s,d) . As mentioned earlier, our upper bound in Theorem 1 does *not* impose a global boundedness condition, whereas in contrast, the analysis of Koltchinskii and Yuan (2010), or KY for short, does impose such a global boundedness condition. Under global boundedness, their work provides rates on the $L^2(\mathbb{P})$ and $L^2(\mathbb{P}_n)$ norm that are of the same order as the results presented here. It is natural to wonder whether or not this difference is actually significant—that is, do the minimax rates for the class of sparse additive models depend on whether or not global boundedness is imposed? In Section 3.5, we answer this question in the affirmative. In particular, Theorem 3 and Corollary 3 provide upper bounds on the minimax rates, as measured in either the $L^2(\mathbb{P})$ and $L^2(\mathbb{P}_n)$ -norms, under a global boundedness condition. These rates are faster than those of Theorem 3 in the KY paper, in particular showing that the KY rates are not optimal for problems with $s = \Omega(\sqrt{n})$. In this way, we see that the assumption of global boundedness, though relatively innocuous for classical (low-dimensional) non-parametric problems, can be quite limiting in high dimensions.

The remainder of the paper is organized as follows. In Section 2, we provide background on kernel spaces and the class of sparse additive models considered in this paper. Section 3 is devoted to the statement of our main results and discussion of their consequences; it includes description of our method, the upper bounds on the convergence rate that it achieves, and a matching set of minimax lower bounds. Section 3.5 investigates the restrictiveness of the global uniform boundedness assumption and in particular, Theorem 3 and Corollary 3 demonstrate that there are classes of globally bounded functions for which faster rates are possible. Section 4 is devoted to the proofs of

our three main theorems, with the more technical details deferred to the Appendices. We conclude with a discussion in Section 5.

2. Background and Problem Set-up

We begin with some background on reproducing kernel Hilbert spaces, before providing a precise definition of the class of sparse additive models studied in this paper.

2.1 Reproducing Kernel Hilbert Spaces

Given a subset $\mathcal{X} \subset \mathbb{R}$ and a probability measure \mathbb{Q} on \mathcal{X} , we consider a Hilbert space $\mathcal{H} \subset L^2(\mathbb{Q})$, meaning a family of functions $g : \mathcal{X} \rightarrow \mathbb{R}$, with $\|g\|_{L^2(\mathbb{Q})} < \infty$, and an associated inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ under which \mathcal{H} is complete. The space \mathcal{H} is a reproducing kernel Hilbert space (RKHS) if there exists a symmetric function $\mathbb{K} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ such that for each $x \in \mathcal{X}$: (a) the function $\mathbb{K}(\cdot, x)$ belongs to the Hilbert space \mathcal{H} , and (b) we have the reproducing relation $f(x) = \langle f, \mathbb{K}(\cdot, x) \rangle_{\mathcal{H}}$ for all $f \in \mathcal{H}$. Any such kernel function must be positive semidefinite; under suitable regularity conditions, Mercer’s theorem (1909) guarantees that the kernel has an eigen-expansion of the form

$$\mathbb{K}(x, x') = \sum_{k=1}^{\infty} \mu_k \phi_k(x) \phi_k(x'), \tag{3}$$

where $\mu_1 \geq \mu_2 \geq \mu_3 \geq \dots \geq 0$ are a non-negative sequence of eigenvalues, and $\{\phi_k\}_{k=1}^{\infty}$ are the associated eigenfunctions, taken to be orthonormal in $L^2(\mathbb{Q})$. The decay rate of these eigenvalues will play a crucial role in our analysis, since they ultimately determine the rate v_n for the univariate RKHS’s in our function classes.

Since the eigenfunctions $\{\phi_k\}_{k=1}^{\infty}$ form an orthonormal basis, any function $f \in \mathcal{H}$ has an expansion of the $f(x) = \sum_{k=1}^{\infty} a_k \phi_k(x)$, where $a_k = \langle f, \phi_k \rangle_{L^2(\mathbb{Q})} = \int_{\mathcal{X}} f(x) \phi_k(x) d\mathbb{Q}(x)$ are (generalized) Fourier coefficients. Associated with any two functions in \mathcal{H} —say $f = \sum_{k=1}^{\infty} a_k \phi_k$ and $g = \sum_{k=1}^{\infty} b_k \phi_k$ —are two distinct inner products. The first is the usual inner product in $L^2(\mathbb{Q})$, $\langle f, g \rangle_{L^2(\mathbb{Q})} := \int_{\mathcal{X}} f(x) g(x) d\mathbb{Q}(x)$. By Parseval’s theorem, it has an equivalent representation in terms of the expansion coefficients—namely

$$\langle f, g \rangle_{L^2(\mathbb{Q})} = \sum_{k=1}^{\infty} a_k b_k.$$

The second inner product, denoted $\langle f, g \rangle_{\mathcal{H}}$, is the one that defines the Hilbert space; it can be written in terms of the kernel eigenvalues and generalized Fourier coefficients as

$$\langle f, g \rangle_{\mathcal{H}} = \sum_{k=1}^{\infty} \frac{a_k b_k}{\mu_k}.$$

Using this definition, the Hilbert ball of unit radius for a kernel with eigenvalues $\{\mu_k\}_{k=1}^{\infty}$ and eigenfunctions $\{\phi_k\}_{k=1}^{\infty}$ is given by

$$\mathbb{B}_{\mathcal{H}}(1) := \left\{ f = \sum_{k=1}^{\infty} a_k \phi_k \mid \sum_{k=1}^{\infty} \frac{a_k^2}{\mu_k} \leq 1 \right\}.$$

For more background on reproducing kernel Hilbert spaces, we refer the reader to various standard references (e.g., Aronszajn, 1950; Saitoh, 1988; Schölkopf and Smola, 2002; Wahba, 1990).

2.2 Sparse Additive Models Over RKHS

For each $j = 1, \dots, d$, let $\mathcal{H}_j \subset L^2(\mathbb{Q})$ be a reproducing kernel Hilbert space of univariate functions on the domain $\mathcal{X} \subset \mathbb{R}$. We assume that

$$\mathbb{E}[f_j(x)] = \int_{\mathcal{X}} f_j(x) d\mathbb{Q}(x) = 0 \quad \text{for all } f_j \in \mathcal{H}_j, \text{ and for each } j = 1, 2, \dots, d.$$

As will be clarified momentarily, our observation model (5) allows for the possibility of a non-zero mean \bar{f} , so that there is no loss of generality in this assumption. For a given subset $S \subset \{1, 2, \dots, d\}$, we define

$$\mathcal{H}(S) := \left\{ f = \sum_{j \in S} f_j \mid f_j \in \mathcal{H}_j, \text{ and } f_j \in \mathbb{B}_{\mathcal{H}_j}(1) \forall j \in S \right\},$$

corresponding to the class of functions $f : \mathcal{X}^d \rightarrow \mathbb{R}$ that decompose as sums of univariate functions on co-ordinates lying within the set S . Note that $\mathcal{H}(S)$ is also (a subset of) a reproducing kernel Hilbert space, in particular with the norm

$$\|f\|_{\mathcal{H}(S)}^2 = \sum_{j \in S} \|f_j\|_{\mathcal{H}_j}^2,$$

where $\|\cdot\|_{\mathcal{H}_j}$ denotes the norm on the univariate Hilbert space \mathcal{H}_j . Finally, for $s \in \{1, 2, \dots, \lfloor d/2 \rfloor\}$, we define the function class

$$\mathcal{F}_{d,s,\mathcal{H}} := \bigcup_{\substack{S \subset \{1,2,\dots,d\} \\ |S|=s}} \mathcal{H}(S). \quad (4)$$

To ease notation, we frequently adopt the shorthand $\mathcal{F} = \mathcal{F}_{d,s,\mathcal{H}}$, but the reader should recall that \mathcal{F} depends on the choice of Hilbert spaces $\{\mathcal{H}_j\}_{j=1}^d$, and moreover, that we are actually studying a *sequence of function classes* indexed by (d, s) .

Now let $\mathbb{P} = \mathbb{Q}^d$ denote the product measure on the space $\mathcal{X}^d \subseteq \mathbb{R}^d$. Given an arbitrary $f^* \in \mathcal{F}$, we consider the observation model

$$y_i = \bar{f} + f^*(x_i) + w_i, \quad \text{for } i = 1, 2, \dots, n, \quad (5)$$

where $\{w_i\}_{i=1}^n$ is an i.i.d. sequence of standard normal variates, and $\{x_i\}_{i=1}^n$ is a sequence of design points in \mathbb{R}^d , sampled in an i.i.d. manner from \mathbb{P} .

Given an estimate \hat{f} , our goal is to bound the error $\hat{f} - f^*$ under two norms. The first is the *usual* $L^2(\mathbb{P})$ norm on the space \mathcal{F} ; given the product structure of \mathbb{P} and the additive nature of any $f \in \mathcal{F}$, it has the additive decomposition $\|f\|_{L^2(\mathbb{P})}^2 = \sum_{j=1}^d \|f_j\|_{L^2(\mathbb{Q})}^2$. In addition, we consider the error in the *empirical* $L^2(\mathbb{P}_n)$ -norm defined by the sample $\{x_i\}_{i=1}^n$, defined as

$$\|f\|_{L^2(\mathbb{P}_n)}^2 := \frac{1}{n} \sum_{i=1}^n f^2(x_i).$$

Unlike the $L^2(\mathbb{P})$ norm, this norm does not decouple across the dimensions, but part of our analysis will establish an approximate form of such decoupling. For shorthand, we frequently use the notation $\|f\|_2 = \|f\|_{L^2(\mathbb{P})}$ and $\|f\|_n = \|f\|_{L^2(\mathbb{P}_n)}$ for a d -variate function $f \in \mathcal{F}$. With a minor abuse of notation, for a univariate function $f_j \in \mathcal{H}_j$, we also use the shorthands $\|f_j\|_2 = \|f_j\|_{L^2(\mathbb{Q})}$ and $\|f_j\|_n = \|f_j\|_{L^2(\mathbb{Q}_n)}$.

3. Main Results and Their Consequences

This section is devoted to the statement of our three main results, and discussion of some of their consequences. We begin in Section 3.1 by describing a regularized M -estimator for sparse additive models, and we state our upper bounds for this estimator in Section 3.2. We illustrate our upper bounds for various concrete instances of kernel classes. In Section 3.3, we state minimax lower bounds on the $L^2(\mathbb{P})$ error over the class $\mathcal{F}_{d,s,\mathcal{H}}$, which establish the optimality of our procedure. In Section 3.4, we provide a detailed comparison between our results to past work, and in Section 3.5 we discuss the effect of global boundedness conditions on optimal rates.

3.1 A Regularized M -Estimator For Sparse Additive Models

For any function of the form $f = \sum_{j=1}^d f_j$, the $(L^2(\mathbb{Q}_n), 1)$ and $(\mathcal{H}, 1)$ -norms are given by

$$\|f\|_{n,1} := \sum_{j=1}^d \|f_j\|_n, \quad \text{and} \quad \|f\|_{\mathcal{H},1} := \sum_{j=1}^d \|f_j\|_{\mathcal{H}},$$

respectively. Using this notation and defining the sample mean $\bar{y}_n = \frac{1}{n} \sum_{i=1}^n y_i$, we define the cost functional

$$\mathcal{L}(f) = \frac{1}{2n} \sum_{i=1}^n (y_i - \bar{y}_n - f(x_i))^2 + \lambda_n \|f\|_{n,1} + \rho_n \|f\|_{\mathcal{H},1}.$$

The cost functional $\mathcal{L}(f)$ is least-squares loss with a sparsity penalty $\|f\|_{n,1}$ and a smoothness penalty $\|f\|_{\mathcal{H},1}$. Here (λ_n, ρ_n) are a pair of positive regularization parameters whose choice will be specified by our theory. Given this cost functional, we then consider the M -estimator

$$\hat{f} \in \arg \min_f \mathcal{L}(f) \quad \text{subject to } f = \sum_{j=1}^d f_j \text{ and } \|f_j\|_{\mathcal{H}} \leq 1 \text{ for all } j = 1, 2, \dots, d. \quad (6)$$

In this formulation (6), the problem is infinite-dimensional in nature, since it involves optimization over Hilbert spaces. However, an attractive feature of this M -estimator is that, as a consequence of the representer theorem (Kimeldorf and Wahba, 1971), it can be reduced to an equivalent convex program in $\mathbb{R}^n \times \mathbb{R}^d$. In particular, for each $j = 1, 2, \dots, d$, let \mathbb{K}^j denote the kernel function for co-ordinate j . Using the notation $x_i = (x_{i1}, x_{i2}, \dots, x_{id})$ for the i^{th} sample, we define the collection of empirical kernel matrices $K^j \in \mathbb{R}^{n \times n}$, with entries $K_{i\ell}^j = \mathbb{K}^j(x_{ij}, x_{\ell j})$. By the representer theorem, any solution \hat{f} to the variational problem (6) can be written in the form

$$\hat{f}(z_1, \dots, z_d) = \sum_{i=1}^n \sum_{j=1}^d \hat{\alpha}_{ij} \mathbb{K}^j(z_j, x_{ij}),$$

for a collection of weights $\{\hat{\alpha}_j \in \mathbb{R}^n, j = 1, \dots, d\}$. The optimal weights $(\hat{\alpha}_1, \dots, \hat{\alpha}_d)$ are any minimizer to the following convex program:

$$\arg \min_{\substack{\alpha_j \in \mathbb{R}^n \\ \alpha_j^T K^j \alpha_j \leq 1}} \left\{ \frac{1}{2n} \|y - \bar{y}_n - \sum_{j=1}^d K^j \alpha_j\|_2^2 + \lambda_n \sum_{j=1}^d \sqrt{\frac{1}{n} \|K^j \alpha_j\|_2^2} + \rho_n \sum_{j=1}^d \sqrt{\alpha_j^T K^j \alpha_j} \right\}. \quad (7)$$

This problem is a second-order cone program (SOCP), and there are various algorithms for finding a solution to arbitrary accuracy in time polynomial in (n, d) , among them interior point methods (e.g., see §11 in Boyd and Vandenberghe 2004).

Various combinations of sparsity and smoothness penalties have been used in past work on sparse additive models. For instance, the method of Ravikumar et al. (2009) is based on least-squares loss regularized with single sparsity constraint, and separate smoothness constraints for each univariate function. They solve the resulting optimization problem using a back-fitting procedure. Koltchinskii and Yuan (2008) develop a method based on least-squares loss combined with a single penalty term $\sum_{j=1}^d \|f_j\|_{\mathcal{H}}$. Their method also leads to an SOCP if \mathcal{H} is a reproducing kernel Hilbert space, but differs from the program (7) in lacking the additional sparsity penalties. Meier et al. (2009) analyzed least-squares regularized with a penalty term of the form $\sum_{j=1}^d \sqrt{\lambda_1 \|f_j\|_n^2 + \lambda_2 \|f_j\|_{\mathcal{H}}^2}$, where λ_1 and λ_2 are a pair of regularization parameters. In their method, λ_1 controls the sparsity while λ_2 controls the smoothness. If \mathcal{H} is an RKHS, the method in Meier et al. (2009) reduces to an ordinary group Lasso problem on a different set of variables, which can be cast as a quadratic program. The more recent work of Koltchinskii and Yuan (2010) is based on essentially the same estimator as problem (6), except that we allow for a non-zero mean for the function, and estimate it as well. In addition, the KY analysis involves a stronger condition of global boundedness. We provide a more in-depth comparison of our analysis and results with the past work listed above in Sections 3.4 and 3.5.

3.2 Upper Bound

We now state a result that provides upper bounds on the estimation error achieved by the estimator (6), or equivalently (7). To simplify presentation, we state our result in the special case that the univariate Hilbert space $\mathcal{H}_j, j = 1, \dots, d$ are all identical, denoted by \mathcal{H} . However, the analysis and results extend in a straightforward manner to the general setting of distinct univariate Hilbert spaces, as we discuss following the statement of Theorem 1.

Let $\mu_1 \geq \mu_2 \geq \dots \geq 0$ denote the non-negative eigenvalues of the kernel operator defining the univariate Hilbert space \mathcal{H} , as defined in Equation (3), and define the function

$$Q_{\sigma,n}(t) := \frac{1}{\sqrt{n}} \left[\sum_{\ell=1}^{\infty} \min\{t^2, \mu_{\ell}\} \right]^{1/2}.$$

Let $v_n > 0$ be the smallest positive solution to the inequality

$$40v_n^2 \geq Q_{\sigma,n}(v_n), \tag{8}$$

where the 40 is simply used for technical convenience. We refer to v_n as the *critical univariate rate*, as it is the minimax-optimal rate for $L^2(\mathbb{P})$ -estimation of a single univariate function in the Hilbert space \mathcal{H} (e.g., Mendelson, 2002; van de Geer, 2000). This quantity will be referred to throughout the remainder of the paper.

Our choices of regularization parameters are specified in terms of the quantity

$$\gamma_n := \kappa \max \left\{ v_n, \sqrt{\frac{\log d}{n}} \right\}, \tag{9}$$

where κ is a fixed constant that we choose later. We assume that each function within the unit ball of the univariate Hilbert space is uniformly bounded by a constant multiple of its Hilbert norm—that is, for each $j = 1, \dots, d$ and each $f_j \in \mathcal{H}$,

$$\|f_j\|_\infty := \sup_{x_j} |f_j(x_j)| \leq c \|f_j\|_{\mathcal{H}}. \tag{10}$$

This condition is satisfied for many kernel classes including Sobolev spaces, and any univariate RKHS in which the kernel function² bounded uniformly by c . Such a condition is routinely imposed for proving upper bounds on rates of convergence for non-parametric least squares in the univariate case $d = 1$ (see, e.g., Stone, 1985; van de Geer, 2000). Note that this univariate boundedness does not imply that the multivariate functions $f = \sum_{j \in S} f_j$ in \mathcal{F} are uniformly bounded independently of (d, s) ; rather, since such functions are the sum of s terms, they can take on values of the order \sqrt{s} .

The following result applies to any class $\mathcal{F}_{d,s,\mathcal{H}}$ of sparse additive models based on a univariate Hilbert space satisfying condition (10), and to the estimator (6) based on n i.i.d. samples $(x_i, y_i)_{i=1}^n$ from the observation model (5).

Theorem 1 *Let \hat{f} be any minimizer of the convex program (6) with regularization parameters $\lambda_n \geq 16\gamma_n$ and $\rho_n \geq 16\gamma_n^2$. Then provided that $n\gamma_n^2 = \Omega(\log(1/\gamma_n))$, there are universal constants (C, c_1, c_2) such that*

$$\mathbb{P} \left[\max \{ \|\hat{f} - f^*\|_2^2, \|\hat{f} - f^*\|_n^2 \} \geq C \{ s\lambda_n^2 + s\rho_n \} \right] \leq c_1 \exp(-c_2 n\gamma_n^2).$$

We provide the proof of Theorem 1 in Section 4.1.

3.2.1 REMARKS

First, the technical condition $n\gamma_n^2 = \Omega(\log(1/\gamma_n))$ is quite mild, and satisfied in most cases of interest, among them the kernels considered below in Corollaries 1 and 2.

Second, note that setting $\lambda_n = c\gamma_n$ and $\rho_n = c\gamma_n^2$ for some constant $c \in [16, \infty)$ yields the rate $\Theta(s\gamma_n^2 + s\rho_n) = \Theta(\frac{s \log d}{n} + sv_n^2)$. This rate may be interpreted as the sum of a subset selection term $(\frac{s \log d}{n})$ and an s -dimensional estimation term (sv_n^2) . Note that the subset selection term $(\frac{s \log d}{n})$ is independent of the choice of Hilbert space \mathcal{H} , whereas the s -dimensional estimation term is independent of the ambient dimension d . Depending on the scaling of the triple (n, d, s) and the smoothness of the univariate RKHS \mathcal{H} , either the subset selection term or function estimation term may dominate. In general, if $\frac{\log d}{n} = o(v_n^2)$, the s -dimensional estimation term dominates, and vice versa otherwise. At the boundary, the scalings of the two terms are equivalent.

Finally, for clarity, we have stated our result in the case where the univariate Hilbert space \mathcal{H} is identical across all co-ordinates. However, our proof extends with only notational changes to the general setting, in which each co-ordinate j is endowed with a (possibly distinct) Hilbert space \mathcal{H}_j . In this case, the M -estimator returns a function \hat{f} such that (with high probability)

$$\max \{ \|\hat{f} - f^*\|_n^2, \|\hat{f} - f^*\|_2^2 \} \leq C \left\{ \frac{s \log d}{n} + \sum_{j \in S} v_{n,j}^2 \right\},$$

2. Indeed, we have

$$\sup_{x_j} |f_j(x_j)| = \sup_{x_j} |\langle f_j(\cdot), \mathbb{K}(\cdot, x_j) \rangle_{\mathcal{H}}| \leq \sup_{x_j} \sqrt{\mathbb{K}(x_j, x_j)} \|f_j\|_{\mathcal{H}}.$$

where $v_{n,j}$ is the critical univariate rate associated with the Hilbert space \mathcal{H}_j , and S is the subset on which f^* is supported.

Theorem 1 has a number of corollaries, obtained by specifying particular choices of kernels. First, we discuss m -rank operators, meaning that the kernel function \mathbb{K} can be expanded in terms of m eigenfunctions. This class includes linear functions, polynomial functions, as well as any function class based on finite dictionary expansions. First we present a corollary for finite-rank kernel classes.

Corollary 1 *Under the same conditions as Theorem 1, consider an univariate kernel with finite rank m . Then any solution \hat{f} to the problem (6) with $\lambda_n = c\gamma_n$ and $\rho_n = c\gamma_n^2$ with $16 \leq c < \infty$ satisfies*

$$\mathbb{P} \left[\max \{ \|\hat{f} - f^*\|_n^2, \|\hat{f} - f^*\|_2^2 \} \geq C \left\{ \frac{s \log d}{n} + s \frac{m}{n} \right\} \right] \leq c_1 \exp(-c_2(m + \log d)).$$

Proof : It suffices to show that the critical univariate rate (8) satisfies the scaling $v_n^2 = O(m/n)$. For a finite-rank kernel and any $t > 0$, we have

$$Q_{\mathcal{G},n}(t) = \frac{1}{\sqrt{n}} \sqrt{\sum_{j=1}^m \min\{t^2, \mu_j\}} \leq t \sqrt{\frac{m}{n}},$$

from which the claim follows by the definition (8). ■

Next, we present a result for the RKHS's with infinitely many eigenvalues, but whose eigenvalues decay at a rate $\mu_k \simeq (1/k)^{2\alpha}$ for some parameter $\alpha > 1/2$. Among other examples, this type of scaling covers the case of Sobolev spaces, say consisting of functions with α derivatives (e.g., Birman and Solomjak, 1967; Gu, 2002).

Corollary 2 *Under the same conditions as Theorem 1, consider an univariate kernel with eigenvalue decay $\mu_k \simeq (1/k)^{2\alpha}$ for some $\alpha > 1/2$. Then the kernel estimator defined in (6) with $\lambda_n = c\gamma_n$ and $\rho_n = c\gamma_n^2$ with $16 \leq c < \infty$ satisfies*

$$\mathbb{P} \left[\max \{ \|\hat{f} - f^*\|_n^2, \|\hat{f} - f^*\|_2^2 \} \geq C \left\{ \frac{s \log d}{n} + s \left(\frac{1}{n} \right)^{\frac{2\alpha}{2\alpha+1}} \right\} \right] \leq c_1 \exp(-c_2(n^{\frac{1}{2\alpha+1}} + \log d)).$$

Proof : As in the previous corollary, we need to compute the critical univariate rate v_n . Given the assumption of polynomial eigenvalue decay, a truncation argument shows that $Q_{\mathcal{G},n}(t) = O\left(t^{\frac{1-\frac{1}{2\alpha}}{\sqrt{n}}}\right)$.

Consequently, the critical univariate rate (8) satisfies the scaling $v_n^2 \asymp v_n^{1-\frac{1}{2\alpha}}/\sqrt{n}$, or equivalently, $v_n^2 \asymp n^{-\frac{2\alpha}{2\alpha+1}}$. ■

3.3 Minimax Lower Bounds

In this section, we derive lower bounds on the minimax error in the $L^2(\mathbb{P})$ -norm that complement the achievability results derived in Theorem 1. Given the function class \mathcal{F} , we define the minimax $L^2(\mathbb{P})$ -error $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})$ to be the largest quantity such that

$$\inf_{\hat{f}_n} \sup_{f^* \in \mathcal{F}} \mathbb{P}_{f^*} [\|\hat{f}_n - f^*\|_2^2 \geq \mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})] \geq 1/2, \quad (11)$$

where the infimum is taken over all measurable functions of the n samples $\{(x_i, y_i)\}_{i=1}^n$, and \mathbb{P}_{f^*} denotes the data distribution when the unknown function is f^* . Given this definition, note that Markov's inequality implies that

$$\inf_{\hat{f}_n} \sup_{f^* \in \mathcal{F}} \mathbb{E} \|\hat{f}_n - f^*\|_2^2 \geq \frac{\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})}{2}.$$

Central to our proof of the lower bounds is the metric entropy structure of the univariate reproducing kernel Hilbert spaces. More precisely, our lower bounds depend on the *packing entropy*, defined as follows. Let (\mathcal{G}, ρ) be a totally bounded metric space, consisting of a set \mathcal{G} and a metric $\rho : \mathcal{G} \times \mathcal{G} \rightarrow \mathbb{R}_+$. An ε -packing of \mathcal{G} is a collection $\{f^1, \dots, f^M\} \subset \mathcal{G}$ such that $\rho(f^i, f^j) \geq \varepsilon$ for all $i \neq j$. The ε -packing number $M(\varepsilon; \mathcal{G}, \rho)$ is the cardinality of the largest ε -packing. The packing entropy is simply the logarithm of the packing number, namely the quantity $\log M(\varepsilon; \mathcal{G}, \rho)$, to which we also refer as the metric entropy. In this paper, we derive explicit minimax lower bounds for two different scalings of the univariate metric entropy.

3.3.1 LOGARITHMIC METRIC ENTROPY

There exists some $m > 0$ such that

$$\log M(\varepsilon; \mathbb{B}_{\mathcal{H}}(1), L^2(\mathbb{P})) \simeq m \log(1/\varepsilon) \quad \text{for all } \varepsilon \in (0, 1). \quad (12)$$

Function classes with metric entropy of this type include linear functions (for which $m = k$), univariate polynomials of degree k (for which $m = k + 1$), and more generally, any function space with finite VC-dimension (van der Vaart and Wellner, 1996). This type of scaling also holds for any RKHS based on a kernel with rank m (e.g., see Carl and Triebel, 1980), and these finite-rank kernels include both linear and polynomial functions as special cases.

3.3.2 POLYNOMIAL METRIC ENTROPY

There exists some $\alpha > 0$ such that

$$\log M(\varepsilon; \mathbb{B}_{\mathcal{H}}(1), L^2(\mathbb{P})) \simeq (1/\varepsilon)^{1/\alpha} \quad \text{for all } \varepsilon \in (0, 1). \quad (13)$$

Various types of Sobolev/Besov classes exhibit this type of metric entropy decay (e.g., Birman and Solomjak, 1967; Gu, 2002). In fact, any RKHS in which the kernel eigenvalues decay at a rate $k^{-2\alpha}$ have a metric entropy with this scaling (Carl and Stephani, 1990; Carl and Triebel, 1980).

We are now equipped to state our lower bounds on the minimax risk (11):

Theorem 2 *Given n i.i.d. samples from the sparse additive model (5) with sparsity $s \leq d/4$, there is an universal constant $C > 0$ such that:*

- (a) *For a univariate class \mathcal{H} with logarithmic metric entropy (12) indexed by parameter m , we have*

$$\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}) \geq C \left\{ \frac{s \log(d/s)}{n} + s \frac{m}{n} \right\}.$$

(b) For a univariate class \mathcal{H} with polynomial metric entropy (13) indexed by α , we have

$$\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}) \geq C \left\{ \frac{s \log(d/s)}{n} + s \left(\frac{1}{n} \right)^{\frac{2\alpha}{2\alpha+1}} \right\}.$$

The proof of Theorem 2 is provided in Section 4.2. The most important consequence of Theorem 2 is in establishing the minimax-optimality of the results given in Corollary 1 and 2; in particular, in the regime of sub-linear sparsity (i.e., for which $\log d = O(\log(d/s))$), the combination of Theorem 2 with these corollaries identifies the minimax rates up to constant factors.

3.4 Comparison With Other Estimators

It is interesting to compare these convergence rates in $L^2(\mathbb{P}_n)$ error with those established in the past work. Ravikumar et al. (2009) show that any solution to their back-fitting method is consistent in terms of mean-squared error risk (see Theorem 3 in their paper), but their analysis does not allow $s \rightarrow \infty$. The method of Koltchinskii and Yuan (2008) is based regularizing the least-squares loss with the $(\mathcal{H}, 1)$ -norm penalty—that is, the regularizer $\sum_{j=1}^d \|f_j\|_{\mathcal{H}}$; Theorem 2 in their paper provides a rate that holds for the triple (n, d, s) tending to infinity. In quantitative terms, however, their rates are looser than those given here; in particular, their bound includes a term of the order $\frac{s^3 \log d}{n}$, which is larger than the bound in Theorem 1. Meier et al. (2009) analyze a different M -estimator to the one we analyze in this paper. Rather than adding two separate $(\mathcal{H}, 1)$ -norm and an $(\|\cdot\|_n, 1)$ -norm penalties, they combine the two terms into a single sparsity and smoothness penalty. For their estimator, Meier et al. (2009) establish a convergence rate of the form $O(s(\frac{\log d}{n})^{\frac{2\alpha}{2\alpha+1}})$ in the case of α -smooth Sobolev spaces (see Theorem 1 in their paper). Note that relative to optimal rates given here in Theorem 2(b), this scaling is sub-optimal: more precisely, we either have $\frac{\log d}{n} < (\frac{\log d}{n})^{\frac{2\alpha}{2\alpha+1}}$, when the subset selection term dominates, or $(\frac{1}{n})^{\frac{2\alpha}{2\alpha+1}} < (\frac{\log d}{n})^{\frac{2\alpha}{2\alpha+1}}$, when the s -dimensional estimation term dominates. In all of the above-mentioned methods, it is unclear whether or not a sharper analysis would yield better rates. Finally, Koltchinskii and Yuan (2010) analyze the same estimator as the M -estimator (6), and for the case of polynomial metric entropy, establish the same rates Theorem 1, albeit under a global boundedness condition. In the following section, we study the implications of this assumption.

3.5 Upper Bounds Under A Global Boundedness Assumption

As discussed previously in the introduction, the paper of Koltchinskii and Yuan (2010), referred to as KY for short, is based on the M -estimator (6). In terms of rates obtained, they establish a convergence rate based on two terms as in Theorem 1, but with a pre-factor that depends on the global quantity

$$B = \sup_{f \in \mathcal{F}_{d,s,\mathcal{H}}} \|f\|_{\infty} = \sup_{f \in \mathcal{F}_{d,s,\mathcal{H}}} \sup_x |f(x)|,$$

assumed to be bounded independently of dimension and sparsity. Such types of global boundedness conditions are fairly standard in classical non-parametric estimation, where they have no effect on minimax rates. In sharp contrast, the analysis of this section shows that for sparse additive models in the regime $s = \Omega(\sqrt{n})$, such global boundedness can *substantially speed up* minimax rates, showing that the rates proven in KY are not minimax optimal for these classes. The underlying insight is as

follows: when the sparsity grows, imposing global boundedness over s -variate functions substantially reduces the effective dimension from its original size s to a lower dimensional quantity, which we denote by $sK_B(s, n)$, and moreover, the quantity $K_B(s, n) \rightarrow 0$ when $s = \Omega(\sqrt{n})$ as described below.

Recall the definition (4) of the function class $\mathcal{F}_{d,s,\mathcal{H}}$. The model considered in the KY paper is the smaller function class

$$\mathcal{F}_{d,s,\mathcal{H}}^*(B) := \bigcup_{\substack{S \subset \{1,2,\dots,d\} \\ |S|=s}} \mathcal{H}(S, B),$$

where $\mathcal{H}(S, B) := \{f = \sum_{j \in S} f_j \mid f_j \in \mathcal{H}, \text{ and } f_j \in \mathbb{B}_{\mathcal{H}}(1) \forall j \in S \text{ and } \|f\|_\infty \leq B\}$.

The following theorem provides sharper rates for the Sobolev case, in which each univariate Hilbert space has eigenvalues decaying as $\mu_k \simeq k^{-2\alpha}$ for some smoothness parameter $\alpha > 1/2$. Our probabilistic bounds involve the quantity

$$\delta_n := \max\left(\sqrt{\frac{s \log(d/s)}{n}}, B\left(\frac{s^{\frac{1}{2}} \log s}{n}\right)^{1/4}\right), \quad (14)$$

and our rates are stated in terms of the function

$$K_B(s, n) := B\sqrt{\log s}(s^{-1/2\alpha}n^{1/(4\alpha+2)})^{2\alpha-1},$$

where it should be noted that $K_B(s, n) \rightarrow 0$ if $s = \Omega(\sqrt{n})$.

With this notation, we have the following *upper bound* on the minimax risk over the function class $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$.

Theorem 3 *Consider any RKHS \mathcal{H} with eigenvalue decay $k^{-2\alpha}$, and uniformly bounded eigenfunctions (i.e., $\|\phi_k\|_\infty \leq C < \infty$ for all k). Then there are universal constants (c_1, c_2, κ) such that with probability greater than $1 - 2\exp(-c_1 n \delta_n^2)$, we have*

$$\min_{\hat{f}} \max_{f^* \in \mathcal{F}_{d,s,\mathcal{H}}^*(B)} \|\hat{f} - f^*\|_2^2 \leq \underbrace{\kappa^2(1+B)Csn^{-\frac{2\alpha}{2\alpha+1}} \left(K_B(s, n) + n^{-1/(2\alpha+1)} \log(d/s)\right)}_{\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}^*(B))}, \quad (15)$$

as long as $n\delta_n^2 = \Omega(\log(1/\delta_n))$.

We provide the proof of Theorem 3 in Section 4.3; it is based on analyzing directly the least-squares estimator over $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$. The assumption that $\|\phi_k\|_\infty \leq C < \infty$ for all k includes the usual Sobolev spaces in which ϕ_k are (rescaled) Fourier basis functions. An immediate consequence of Theorem 3 is that the minimax rates over the function class $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$ can be strictly faster than minimax rates for the class $\mathcal{F}_{d,s,\mathcal{H}}$, which does not impose global boundedness. Recall that the minimax lower bound from Theorem 2 (b) is based on the quantity

$$\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}) := C_1 \left\{ s \left(\frac{1}{n}\right)^{\frac{2\alpha}{2\alpha+1}} + \frac{s \log(d/s)}{n} \right\} = C_1 s n^{-\frac{2\alpha}{2\alpha+1}} \left(1 + n^{-1/(2\alpha+1)} \log(d/s) \right),$$

for a universal constant C_1 . Note that up to constant factors, the achievable rate (15) from Theorem 3 is the same except that the term 1 is replaced by the function $K_B(s, n)$. Consequently, for scalings of (s, n) such that $K_B(s, n) \rightarrow 0$, global boundedness conditions lead to strictly faster rates.

Corollary 3 *Under the conditions of Theorem 3, we have*

$$\frac{\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})}{\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}^*(B))} \geq \frac{C_1(1+n^{-1/(2\alpha+1)}\log(d/s))}{C\kappa^2(1+B)(K_B(s,n)+n^{-1/(2\alpha+1)}\log(d/s))} \rightarrow +\infty$$

whenever $B = O(1)$ and $K_B(s,n) \rightarrow 0$.

3.5.1 REMARKS

The quantity $K_B(s,n)$ is guaranteed to decay to zero as long as the sparsity index s grows in a non-trivial way with the sample size. For instance, if we have $s = \Omega(\sqrt{n})$ for a problem of dimension $d = O(n^\beta)$ for any $\beta \geq 1/2$, then it can be verified that $K_B(s,n) = o(1)$. As an alternative view of the differences, it can be noted that there are scalings of (n,s,d) for which the minimax rate $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})$ over $\mathcal{F}_{d,s,\mathcal{H}}$ is constant—that is, does not vanish as $n \rightarrow +\infty$ —while the minimax rate $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}^*(B))$ does vanish. As an example, consider the Sobolev class with smoothness $\alpha = 2$, corresponding to twice-differentiable functions. For a sparsity index $s = \Theta(n^{4/5})$, then Theorem 2(b) implies that $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}) = \Omega(1)$, so that the minimax rate over $\mathcal{F}_{d,s,\mathcal{H}}$ is strictly bounded away from zero for all sample sizes. In contrast, under a global boundedness condition, Theorem 3 shows that the minimax rate is upper bounded as $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}}^*(B)) = O(n^{-1/5}\sqrt{\log n})$, which tends to zero.

In summary, Theorem 3 and Theorem 2(b) together show that the minimax rates over $\mathcal{F}_{d,s,\mathcal{H}}$ and $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$ can be drastically different. Thus, global boundedness is a stringent condition in the high-dimensional setting; in particular, the rates given in Theorem 3 of Koltchinskii and Yuan (2010) are not minimax optimal when $s = \Omega(\sqrt{n})$.

4. Proofs

In this section, we provide the proofs of our three main theorems. For clarity in presentation, we split the proofs up into a series of lemmas, with the bulk of the more technical arguments deferred to the appendices. This splitting allows our presentation in Section 4 to be relatively streamlined.

4.1 Proof of Theorem 1

At a high-level, Theorem 1 is based on an appropriate adaptation to the non-parametric setting of various techniques that have been developed for sparse linear regression (e.g., Bickel et al., 2009; Negahban et al., 2009). In contrast to the parametric setting where classical tail bounds are sufficient, controlling the error terms in the non-parametric case requires more advanced techniques from empirical process theory. In particular, we make use of various concentration theorems for Gaussian and empirical processes (e.g., Ledoux, 2001; Massart, 2000; Pisier, 1989; van de Geer, 2000), as well as results on the Rademacher complexity of kernel classes (Bartlett et al., 2005; Mendelson, 2002).

At the core of the proof are three technical lemmas. First, Lemma 1 provides an upper bound on the Gaussian complexity of any function of the form $f = \sum_{j=1}^d f_j$ in terms of the norms $\|\cdot\|_{\mathcal{H},1}$ and $\|\cdot\|_{n,1}$ previously defined. Lemma 2 exploits the notion of decomposability (Negahban et al., 2009), as applied to these norms, in order to show that the error function belongs to a particular cone-shaped set. Finally, Lemma 3 establishes an upper bound on the $L^2(\mathbb{P})$ error of our estimator in terms of the $L^2(\mathbb{P}_n)$ error. The latter lemma can be interpreted as proving that our problem satisfies non-parametric analog of a restricted eigenvalue condition (Bickel et al., 2009), or more

generally, of a restricted strong convexity condition (Negahban et al., 2009). The proof of Lemma 3 involves a new approach that combines the Sudakov minoration (Pisier, 1989) with a one-sided tail bound for non-negative random variables (Chung and Lu, 2006; Einmahl and Mason, 1996).

Throughout the proof, we use C and c_i , $i = 1, 2, 3, 4$ to denote universal constants, independent of (n, d, s) . Note that the precise numerical values of these constants may change from line to line. The reader should recall the definitions of v_n and γ_n from Equations (8) and (9) respectively. For a subset $A \subseteq \{1, 2, \dots, d\}$ and a function of the form $f = \sum_{j=1}^d f_j$, we adopt the convenient notation

$$\|f_A\|_{n,1} := \sum_{j \in A} \|f_j\|_n, \quad \text{and} \quad \|f_A\|_{\mathcal{H},1} := \sum_{j \in A} \|f_j\|_{\mathcal{H}}. \quad (16)$$

We begin by establishing an inequality on the error function $\widehat{\Delta} := \widehat{f} - f^*$. Since \widehat{f} and f^* are, respectively, optimal and feasible for the problem (6), we are guaranteed that $\mathcal{L}(\widehat{f}) \leq \mathcal{L}(f^*)$, and hence that the error function $\widehat{\Delta}$ satisfies the bound

$$\frac{1}{2n} \sum_{i=1}^n (w_i + \bar{f} - \bar{y}_n - \widehat{\Delta}(x_i))^2 + \lambda_n \|\widehat{f}\|_{n,1} + \rho_n \|\widehat{f}\|_{\mathcal{H},1} \leq \frac{1}{2n} \sum_{i=1}^n (w_i + \bar{f} - \bar{y}_n)^2 + \lambda_n \|f^*\|_{n,1} + \rho_n \|f^*\|_{\mathcal{H},1}.$$

Some simple algebra yields the bound

$$\frac{1}{2} \|\widehat{\Delta}\|_n^2 \leq \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| + |\bar{y}_n - \bar{f}| \left| \frac{1}{n} \sum_{i=1}^n \widehat{\Delta}(x_i) \right| + \lambda_n \|\widehat{\Delta}\|_{n,1} + \rho_n \|\widehat{\Delta}\|_{\mathcal{H},1}. \quad (17)$$

Following the terminology of van de Geer (2000), we refer to this bound as our *basic inequality*.

4.1.1 CONTROLLING DEVIATION FROM THE MEAN

Our next step is to control the error due to estimating the mean $|\bar{y}_n - \bar{f}|$. We begin by observing that this error term can be written as $\bar{y}_n - \bar{f} = \frac{1}{n} \sum_{i=1}^n (y_i - \bar{f})$. Next we observe that $y_i - \bar{f} = \sum_{j \in S} f_j^*(x_{ij}) + w_i$ is the sum of the s independent random variables $f_j^*(x_{ij})$, each bounded in absolute value by one, along with the independent sub-Gaussian noise term w_i ; consequently, the variable $y_i - \bar{f}$ is sub-Gaussian with parameter at most $\sqrt{s+1}$. (See, for instance, Lemma 1.4 in Buldygin and Kozachenko 2000). By applying standard sub-Gaussian tail bounds, we have $\mathbb{P}(|\bar{y}_n - \bar{f}| > t) \leq 2 \exp(-\frac{mt^2}{2(s+1)})$, and hence, if we define the event $\mathcal{C}(\gamma_n) = \{|\bar{y}_n - \bar{f}| \leq \sqrt{s}\gamma_n\}$, we are guaranteed

$$\mathbb{P}[\mathcal{C}(\gamma_n)] \geq 1 - 2 \exp(-\frac{n\gamma_n^2}{4}).$$

For the remainder of the proof, we condition on the event $\mathcal{C}(\gamma_n)$. Under this conditioning, the bound (17) simplifies to:

$$\frac{1}{2} \|\widehat{\Delta}\|_n^2 \leq \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| + \sqrt{s}\gamma_n \|\widehat{\Delta}\|_n + \lambda_n \|\widehat{\Delta}\|_{n,1} + \rho_n \|\widehat{\Delta}\|_{\mathcal{H},1},$$

where we have applied the Cauchy-Schwarz inequality to write $\left| \frac{1}{n} \sum_{i=1}^n \widehat{\Delta}(x_i) \right| \leq \|\widehat{\Delta}\|_n$.

4.1.2 CONTROLLING THE GAUSSIAN COMPLEXITY TERM

The following lemma provides control the Gaussian complexity term on the right-hand side of inequality (17) by bounding the Gaussian complexity for the univariate functions $\widehat{\Delta}_j$, $j = 1, 2, \dots, d$ in terms of their $\|\cdot\|_n$ and $\|\cdot\|_{\mathcal{H}}$ norms. In particular, recalling that $\gamma_n = \kappa \max\{\sqrt{\frac{\log d}{n}}, \nu_n\}$, we have the following lemma.

Lemma 1 *Define the event*

$$\mathcal{T}(\gamma_n) := \left\{ \forall j = 1, 2, \dots, d, \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}_j(x_{ij}) \right| \leq 8\gamma_n^2 \|\widehat{\Delta}_j\|_{\mathcal{H}} + 8\gamma_n \|\widehat{\Delta}_j\|_n \right\}.$$

Then under the condition $n\gamma_n^2 = \Omega(\log(1/\gamma_n))$, we have

$$\mathbb{P}(\mathcal{T}(\gamma_n)) \geq 1 - c_1 \exp(-c_2 n \gamma_n^2).$$

The proof of this lemma, provided in Appendix B, uses concentration of measure for Lipschitz functions of Gaussian random variables (e.g., Ledoux, 2001), combined with peeling and weighting arguments from empirical process theory (Alexander, 1987; van de Geer, 2000). In particular, the subset selection term $(\frac{s \log d}{n})$ in Theorem 1 arises from taking the maximum over all d components.

The remainder of our analysis involves conditioning on the event $\mathcal{T}(\gamma_n) \cap \mathcal{C}(\gamma_n)$. Using Lemma 1, when conditioned on the event $\mathcal{T}(\gamma_n) \cap \mathcal{C}(\gamma_n)$ we have:

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\gamma_n \|\widehat{\Delta}\|_n + (16\gamma_n + 2\lambda_n) \|\widehat{\Delta}\|_{n,1} + (16\gamma_n^2 + 2\rho_n) \|\widehat{\Delta}\|_{\mathcal{H},1}. \quad (18)$$

4.1.3 EXPLOITING DECOMPOSABILITY

Recall that S denotes the true support of the unknown function f^* . By the definition (16), we can write $\|\widehat{\Delta}\|_{n,1} = \|\widehat{\Delta}_S\|_{n,1} + \|\widehat{\Delta}_{S^c}\|_{n,1}$, where $\widehat{\Delta}_S := \sum_{j \in S} \widehat{\Delta}_j$ and $\widehat{\Delta}_{S^c} := \sum_{j \in S^c} \widehat{\Delta}_j$. Similarly, we have an analogous representation for $\|\widehat{\Delta}\|_{\mathcal{H},1}$. The next lemma shows that conditioned on the event $\mathcal{T}(\gamma_n)$, the quantities $\|\widehat{\Delta}\|_{\mathcal{H},1}$ and $\|\widehat{\Delta}\|_{n,1}$ are not significantly larger than the corresponding norms as applied to the function $\widehat{\Delta}_S$.

Lemma 2 *Conditioned on the events $\mathcal{T}(\gamma_n)$ and $\mathcal{C}(\gamma_n)$, and with the choices $\lambda_n \geq 16\gamma_n$ and $\rho_n \geq 16\gamma_n^2$, we have*

$$\lambda_n \|\widehat{\Delta}\|_{n,1} + \rho_n \|\widehat{\Delta}\|_{\mathcal{H},1} \leq 4\lambda_n \|\widehat{\Delta}_S\|_{n,1} + 4\rho_n \|\widehat{\Delta}_S\|_{\mathcal{H},1} + \frac{1}{2} s \gamma_n^2. \quad (19)$$

The proof of this lemma, provided in Appendix C, is based on the decomposability (see Negahban et al. 2009) of the $\|\cdot\|_{\mathcal{H},1}$ and $\|\cdot\|_{n,1}$ norms. This lemma allows us to exploit the sparsity assumption, since in conjunction with Lemma 1, we have now bounded the right-hand side of the inequality (18) by terms involving only $\widehat{\Delta}_S$.

For the remainder of the proof of Theorem 1, we assume $\lambda_n \geq 16\gamma_n$ and $\rho_n \geq 16\gamma_n^2$. In particular, still conditioning on $\mathcal{C}(\gamma_n) \cap \mathcal{T}(\gamma_n)$ and applying Lemma 2 to inequality (18), we obtain

$$\begin{aligned} \|\widehat{\Delta}\|_n^2 &\leq 2\sqrt{s}\gamma_n \|\widehat{\Delta}\|_n + 3\lambda_n \|\widehat{\Delta}\|_{n,1} + 3\rho_n \|\widehat{\Delta}\|_{\mathcal{H},1} \\ &\leq 2\sqrt{s}\lambda_n \|\widehat{\Delta}\|_n + 12\lambda_n \|\widehat{\Delta}_S\|_{n,1} + 12\rho_n \|\widehat{\Delta}_S\|_{\mathcal{H},1} + \frac{3}{32} s \rho_n, \end{aligned}$$

Finally, since both \widehat{f}_j and f_j^* belong to $\mathbb{B}_{\mathcal{H}}(1)$, we have $\|\widehat{\Delta}_j\|_{\mathcal{H}} \leq \|\widehat{f}_j\|_{\mathcal{H}} + \|f_j^*\|_{\mathcal{H}} \leq 2$, which implies that $\|\widehat{\Delta}_S\|_{\mathcal{H},1} \leq 2s$, and hence

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\lambda_n\|\widehat{\Delta}\|_n + 12\lambda_n\|\widehat{\Delta}_S\|_{n,1} + 25s\rho_n. \quad (20)$$

4.1.4 UPPER BOUNDING $\|\widehat{\Delta}_S\|_{n,1}$

The next step is to control the term $\|\widehat{\Delta}_S\|_{n,1} = \sum_{j \in S} \|\widehat{\Delta}_j\|_n$ that appears in the upper bound (20). Ideally, we would like to upper bound it by a quantity of the order $\sqrt{s}\|\widehat{\Delta}_S\|_2 = \sqrt{s}\sqrt{\sum_{j \in S} \|\widehat{\Delta}_j\|_2^2}$. Such an upper bound would follow immediately if it were phrased in terms of the population $\|\cdot\|_2$ -norm rather than the empirical- $\|\cdot\|_n$ norm, but there are additional cross-terms with the empirical norm. Accordingly, a somewhat more delicate argument is required, which we provide here. First define the events

$$\mathcal{A}_j(\lambda_n) := \{\|\widehat{\Delta}_j\|_n \leq 2\|\widehat{\Delta}_j\|_2 + \lambda_n\},$$

and $\mathcal{A}(\lambda_n) = \bigcap_{j=1}^d \mathcal{A}_j(\lambda_n)$. By applying Lemma 7 from Appendix A with $t = \lambda_n \geq 16\gamma_n$ and $b = 2$, we conclude that $\|\widehat{\Delta}_j\|_n \leq 2\|\widehat{\Delta}_j\|_2 + \lambda_n$ with probability greater than $1 - c_1 \exp(-c_2 n \lambda_n^2)$. Consequently, if we define the event $\mathcal{A}(\lambda_n) = \bigcap_{j \in S} \mathcal{A}_j(\lambda_n)$, then this tail bound together with the union bound implies that

$$\mathbb{P}[\mathcal{A}^c(\lambda_n)] \leq s c_1 \exp(-c_2 n \lambda_n^2) \leq c_1 \exp(-c'_2 n \lambda_n^2), \quad (21)$$

where we have used the fact that $\lambda_n = \Omega(\sqrt{\frac{\log s}{n}})$. Now, conditioned on the event $\mathcal{A}(\lambda_n)$, we have

$$\begin{aligned} \|\widehat{\Delta}_S\|_{n,1} = \sum_{j \in S} \|\widehat{\Delta}_j\|_n &\leq 2 \sum_{j \in S} \|\widehat{\Delta}_j\|_2 + s\lambda_n \\ &\leq 2\sqrt{s}\|\widehat{\Delta}_S\|_2 + s\lambda_n \leq 2\sqrt{s}\|\widehat{\Delta}\|_2 + s\lambda_n. \end{aligned} \quad (22)$$

Substituting this upper bound (22) on $\|\widehat{\Delta}_S\|_{n,1}$ into our earlier inequality (20) yields

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\lambda_n\|\widehat{\Delta}\|_n + 24\sqrt{s}\lambda_n\|\widehat{\Delta}\|_2 + 12s\lambda_n^2 + 25s\rho_n. \quad (23)$$

At this point, we encounter a challenge due to the unbounded nature of our function class. In particular, if $\|\widehat{\Delta}\|_2$ were upper bounded by $C \max(\|\widehat{\Delta}\|_n, \sqrt{s}\lambda_n, \sqrt{s\rho_n})$, then the upper bound (23) would immediately imply the claim of Theorem 1. If one were to assume global boundedness of the multivariate functions \widehat{f} and f^* , as done in past work of Koltchinskii and Yuan (2010), then an upper bound on $\|\widehat{\Delta}\|_2$ of this form would directly follow from known results (e.g., Theorem 2.1 in Bartlett et al. 2005.) However, since we do not impose global boundedness, we need to develop a novel approach to this final hurdle.

4.1.5 CONTROLLING $\|\widehat{\Delta}\|_2$ FOR UNBOUNDED CLASSES

For the remainder of the proof, we condition on the event $\mathcal{A}(\lambda_n) \cap \mathcal{T}(\gamma_n) \cap \mathcal{C}(\gamma_n)$. We split our analysis into three cases. Throughout the proof, we make use of the quantity

$$\widetilde{\delta}_n := B \max(\sqrt{s}\lambda_n, \sqrt{s\rho_n}), \quad (24)$$

where $B \in (1, \infty)$ is a constant to be chosen later in the argument.

Case 1: If $\|\widehat{\Delta}\|_2 < \|\widehat{\Delta}\|_n$, then combined with inequality (23), we conclude that

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\lambda_n \|\widehat{\Delta}\|_n + 24\sqrt{s}\lambda_n \|\widehat{\Delta}\|_n + 12s\lambda_n^2 + 25s\rho_n.$$

This is a quadratic inequality in terms of the quantity $\|\widehat{\Delta}\|_n$, and some algebra shows that it implies the bound $\|\widehat{\Delta}\|_n \leq 15 \max(\sqrt{s}\lambda_n, \sqrt{s\rho_n})$. By assumption, we then have $\|\widehat{\Delta}\|_2 \leq 15 \max(\sqrt{s}\lambda_n, \sqrt{s\rho_n})$ as well, thereby completing the proof of Theorem 1.

Case 2: If $\|\widehat{\Delta}\|_2 < \tilde{\delta}_n$, then together with the bound (23), we conclude that

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\lambda_n \|\widehat{\Delta}\|_n + 24\sqrt{s}\lambda_n \tilde{\delta}_n + 12s\lambda_n^2 + 25s\rho_n.$$

This inequality is again a quadratic in $\|\widehat{\Delta}\|_n$; moreover, note that by definition (24) of $\tilde{\delta}_n$, we have $s\lambda_n^2 + s\rho_n = O(\tilde{\delta}_n^2)$. Consequently, this inequality implies that $\|\widehat{\Delta}\|_n \leq C\tilde{\delta}_n$ for some constant C . Our starting assumption implies that $\|\widehat{\Delta}\|_2 \leq \tilde{\delta}_n$, so that the claim of Theorem 1 follows in this case.

Case 3: Otherwise, we may assume that $\|\widehat{\Delta}\|_2 \geq \tilde{\delta}_n$ and $\|\widehat{\Delta}\|_2 \geq \|\widehat{\Delta}\|_n$. In this case, the inequality (23) together with the bound $\|\widehat{\Delta}\|_2 \geq \|\widehat{\Delta}\|_n$ implies that

$$\|\widehat{\Delta}\|_n^2 \leq 2\sqrt{s}\lambda_n \|\widehat{\Delta}\|_2 + 24\sqrt{s}\lambda_n \|\widehat{\Delta}\|_2 + 12s\lambda_n^2 + 25s\rho_n. \quad (25)$$

Our goal is to establish a lower bound on the left-hand-side—namely, the quantity $\|\widehat{\Delta}\|_n^2$ —in terms of $\|\widehat{\Delta}\|_2^2$. In order to do so, we consider the function class $\mathcal{G}(\lambda_n, \rho_n)$ defined by functions of the form $g = \sum_{j=1}^d g_j$, and such that

$$\lambda_n \|g\|_{n,1} + \rho_n \|g\|_{\mathcal{H},1} \leq 4\lambda_n \|g_S\|_{n,1} + 4\rho_n \|g_S\|_{\mathcal{H},1} + \frac{1}{32}s\rho_n, \quad (26)$$

$$\|g_S\|_{1,n} \leq 2\sqrt{s} \|g_S\|_2 + s\lambda_n \quad \text{and} \quad (27)$$

$$\|g\|_n \leq \|g\|_2. \quad (28)$$

Conditioned on the events $\mathcal{A}(\gamma_n)$, $\mathcal{T}(\gamma_n)$ and $\mathcal{C}(\gamma_n)$, and with our choices of regularization parameter, we are guaranteed that the error function $\widehat{\Delta}$ satisfies all three of these constraints, and hence that $\widehat{\Delta} \in \mathcal{G}(\lambda_n, \rho_n)$. Consequently, it suffices to establish a lower bound on $\|g\|_n$ that holds uniformly over the class $\mathcal{G}(\lambda_n, \rho_n)$. In particular, define the event

$$\mathcal{B}(\lambda_n, \rho_n) := \left\{ \|g\|_n^2 \geq \|g\|_2^2/2 \quad \text{for all } g \in \mathcal{G}(\lambda_n, \rho_n) \quad \text{such that} \quad \|g\|_2 \geq \tilde{\delta}_n \right\}.$$

The following lemma shows that this event holds with high probability.

Lemma 3 *Under the conditions of Theorem 1, there are universal constants c_i such that*

$$\mathbb{P}[\mathcal{B}(\lambda_n, \rho_n)] \geq 1 - c_1 \exp(-c_2 n \gamma_n^2).$$

We note that this lemma can be interpreted as guaranteeing a version of restricted strong convexity (see Negahban et al., 2009) for the least-squares loss function, suitably adapted to the non-parametric setting. Since we do not assume global boundedness, the proof of this lemma requires a novel technical argument, one which combines a one-sided tail bound for non-negative random

variables (Chung and Lu, 2006; Einmahl and Mason, 1996) with the Sudakov minoration (Pisier, 1989) for the Gaussian complexity. We refer the reader to Appendix D for the details of the proof.

Using Lemma 3 and conditioning on the event $\mathcal{B}(\lambda_n, \rho_n)$, we are guaranteed that $\|\widehat{\Delta}\|_n^2 \geq \|\widehat{\Delta}\|_2^2/2$, and hence, combined with our earlier bound (25), we conclude that

$$\|\widehat{\Delta}\|_2^2 \leq 4\sqrt{s}\lambda_n\|\widehat{\Delta}\|_2 + 48\sqrt{s}\lambda_n\|\widehat{\Delta}\|_2 + 24s\lambda_n^2 + 50s\rho_n.$$

Hence $\|\widehat{\Delta}\|_n \leq \|\widehat{\Delta}\|_2 \leq C \max(\sqrt{s}\lambda_n, \sqrt{s\rho_n})$, completing the proof of the claim in the third case.

In summary, the entire proof is based on conditioning on the three events $\mathcal{T}(\gamma_n)$, $\mathcal{A}(\lambda_n)$ and $\mathcal{B}(\lambda_n, \rho_n)$. From the bound (21) as well as Lemmas 1 and 3, we have

$$\mathbb{P}[\mathcal{T}(\gamma_n) \cap \mathcal{A}(\lambda_n) \cap \mathcal{B}(\lambda_n, \rho_n) \cap \mathcal{C}(\gamma_n)] \geq 1 - c_1 \exp(-c_2 n \gamma_n^2),$$

thereby showing that $\max\{\|\widehat{f} - f^*\|_n^2, \|\widehat{f} - f^*\|_2^2\} \leq C \max(s\lambda_n^2, s\rho_n)$ with the claimed probability. This completes the proof of Theorem 1.

4.2 Proof of Theorem 2

We now turn to the proof of the minimax lower bounds stated in Theorem 2. For both parts (a) and (b), the first step is to follow a standard reduction to testing (see, e.g., Has'minskii, 1978; Yang and Barron, 1999; Yu, 1996) so as to obtain a lower bound on the minimax error $\mathfrak{M}_{\mathbb{P}}(\mathcal{F}_{d,s,\mathcal{H}})$ in terms of the probability of error in a multi-way hypothesis testing. We then apply different forms of the Fano inequality (see Yang and Barron, 1999; Yu, 1996) in order to lower bound the probability of error in this testing problem. Obtaining useful bounds requires a precise characterization of the metric entropy structure of $\mathcal{F}_{d,s,\mathcal{H}}$, as stated in Lemma 4.

4.2.1 REDUCTION TO TESTING

We begin with the reduction to a testing problem. Let $\{f^1, \dots, f^M\}$ be a δ_n -packing of \mathcal{F} in the $\|\cdot\|_2$ -norm, and let Θ be a random variable uniformly distributed over the index set $[M] := \{1, 2, \dots, M\}$. Note that we are using M as a shorthand for the packing number $M(\delta_n; \mathcal{F}, \|\cdot\|_2)$. A standard argument (e.g., Has'minskii, 1978; Yang and Barron, 1999; Yu, 1996) then yields the lower bound

$$\inf_{\widehat{f}} \sup_{f^* \in \mathcal{F}} \mathbb{P}[\|\widehat{f} - f^*\|_2^2 \geq \delta_n^2/2] \geq \inf_{\Theta} \mathbb{P}[\widehat{\Theta} \neq \Theta],$$

where the infimum on the right-hand side is taken over all estimators $\widehat{\Theta}$ that are measurable functions of the data, and take values in the index set $[M]$.

Note that $\mathbb{P}[\widehat{\Theta} \neq \Theta]$ corresponds to the error probability in a multi-way hypothesis test, where the probability is taken over the random choice of Θ , the randomness of the design points $X_1^n := \{x_i\}_{i=1}^n$, and the randomness of the observations $Y_1^n := \{y_i\}_{i=1}^n$. Our initial analysis is performed conditionally on the design points, so that the only remaining randomness in the observations Y_1^n comes from the observation noise $\{w_i\}_{i=1}^n$. From Fano's inequality (Cover and Thomas, 1991), for any estimator $\widehat{\Theta}$, we have $\mathbb{P}[\widehat{\Theta} \neq \Theta \mid X_1^n] \geq 1 - \frac{I_{X_1^n}(\Theta; Y_1^n) + \log 2}{\log M}$, where $I_{X_1^n}(\Theta; Y_1^n)$ denotes the mutual information between Θ and Y_1^n with X_1^n fixed. Taking expectations over X_1^n , we obtain the lower bound

$$\mathbb{P}[\widehat{\Theta} \neq \Theta] \geq 1 - \frac{\mathbb{E}_{X_1^n} [I_{X_1^n}(\Theta; Y_1^n)] + \log 2}{\log M}. \quad (29)$$

The remainder of the proof consists of constructing appropriate packing sets of \mathcal{F} , and obtaining good upper bounds on the mutual information term in the lower bound (29).

4.2.2 CONSTRUCTING APPROPRIATE PACKINGS

We begin with results on packing numbers. Recall that $\log M(\delta; \mathcal{F}, \|\cdot\|_2)$ denotes the δ -packing entropy of \mathcal{F} in the $\|\cdot\|_2$ norm.

Lemma 4 (a) *For all $\delta \in (0, 1)$ and $s \leq d/4$, we have*

$$\log M(\delta; \mathcal{F}, \|\cdot\|_2) = O\left(s \log M\left(\frac{\delta}{\sqrt{s}}; \mathbb{B}_{\mathcal{H}}(1), \|\cdot\|_2\right) + s \log \frac{d}{s}\right).$$

(b) *For a Hilbert class with logarithmic metric entropy (12) and such that $\|f\|_2 \leq \|f\|_{\mathcal{H}}$, there exists set $\{f^1, \dots, f^M\}$ with $\log M \geq C\{s \log(d/s) + sm\}$, and*

$$\delta \leq \|f^k - f^\ell\|_2 \leq 8\delta \quad \text{for all } k \neq \ell \in \{1, 2, \dots, M\}.$$

The proof, provided in Appendix E, is combinatorial in nature. We now turn to the proofs of parts (a) and (b) of Theorem 2.

4.2.3 PROOF OF THEOREM 2(A)

In order to prove this claim, it remains to exploit Lemma 4 in an appropriate way, and to upper bound the resulting mutual information. For the latter step, we make use of the generalized Fano approach (e.g., Yu, 1996).

From Lemma 4, we can find a set $\{f^1, \dots, f^M\}$ that is a δ -packing of \mathcal{F} in ℓ_2 -norm, and such that $\|f^k - f^\ell\|_2 \leq 8\delta$ for all $k, \ell \in [M]$. For $k = 1, \dots, M$, let \mathbb{Q}^k denote the conditional distribution of Y_1^n conditioned on X_1^n and the event $\{\Theta = k\}$, and let $D(\mathbb{Q}^k \|\mathbb{Q}^\ell)$ denote the Kullback-Leibler divergence. From the convexity of mutual information (Cover and Thomas, 1991), we have the upper bound $I_{X_1^n}(\Theta; Y_1^n) \leq \frac{1}{\binom{M}{2}} \sum_{k, \ell=1}^M D(\mathbb{Q}^k \|\mathbb{Q}^\ell)$. Given our linear observation model (5), we have

$$D(\mathbb{Q}^k \|\mathbb{Q}^\ell) = \frac{1}{2\sigma^2} \sum_{i=1}^n (f^k(x_i) - f^\ell(x_i))^2 = \frac{n \|f^k - f^\ell\|_n^2}{2},$$

and hence

$$\mathbb{E}_{X_1^n} [I_{X_1^n}(Y_1^n; \Theta)] \leq \frac{n}{2} \frac{1}{\binom{M}{2}} \sum_{k \neq \ell} \mathbb{E}_{X_1^n} [\|f^k - f^\ell\|_n^2] = \frac{n}{2} \frac{1}{\binom{M}{2}} \sum_{k \neq \ell} \|f^k - f^\ell\|_2^2.$$

Since our packing satisfies $\|f^k - f^\ell\|_2^2 \leq 64\delta^2$, we conclude that

$$\mathbb{E}_{X_1^n} [I_{X_1^n}(Y_1^n; \Theta)] \leq 32n\delta^2.$$

From the Fano bound (29), for any $\delta > 0$ such that $\frac{32n\delta^2 + \log 2}{\log M} < \frac{1}{4}$, then we are guaranteed that $\mathbb{P}[\widehat{\Theta} \neq \Theta] \geq \frac{3}{4}$. From Lemma 4(b), our packing set satisfies $\log M \geq C\{sm + s \log(d/s)\}$, so that so that the choice $\delta^2 = C' \left\{ \frac{sm}{n} + \frac{s \log(d/s)}{n} \right\}$, for a suitably small $C' > 0$, can be used to guarantee the error bound $\mathbb{P}[\widehat{\Theta} \neq \Theta] \geq \frac{3}{4}$.

4.2.4 PROOF OF THEOREM 2(B)

In this case, we use an upper bounding technique due to Yang and Barron (1999) in order to upper bound the mutual information. Although the argument is essentially the same, it does not follow verbatim from their claims—in particular, there are some slight differences due to our initial conditioning—so that we provide the details here. By definition of the mutual information, we have

$$I_{X_1^n}(\Theta; Y_1^n) = \frac{1}{M} \sum_{k=1}^M D(\mathbb{Q}^k \parallel \mathbb{P}_Y),$$

where \mathbb{Q}^k denotes the conditional distribution of Y_1^n given $\Theta = k$ and still with X_1^n fixed, whereas \mathbb{P}_Y denotes the marginal distribution of \mathbb{P}_Y .

Let us define the notion of a covering number, in particular for a totally bounded metric space (\mathcal{G}, ρ) , consisting of a set \mathcal{G} and a metric $\rho : \mathcal{G} \times \mathcal{G} \rightarrow \mathbb{R}_+$. An ε -covering set of \mathcal{G} is a collection $\{f^1, \dots, f^N\}$ of functions such that for all $f \in \mathcal{G}$ there exists $k \in \{1, 2, \dots, N\}$ such that $\rho(f, f^k) \leq \varepsilon$. The ε -covering number $N(\varepsilon; \mathcal{G}, \rho)$ is the cardinality of the smallest ε -covering set.

Now let $\{g^1, \dots, g^N\}$ be an ε -cover of \mathcal{F} in the $\|\cdot\|_2$ norm, for a tolerance ε to be chosen. As argued in Yang and Barron (1999), we have

$$I_{X_1^n}(\Theta; Y_1^n) = \frac{1}{M} \sum_{j=1}^M D(\mathbb{Q}^j \parallel \mathbb{P}_Y) \leq D(\mathbb{Q}^k \parallel \frac{1}{N} \sum_{k=1}^N \mathbb{P}^k),$$

where \mathbb{P}^ℓ denotes the conditional distribution of Y_1^n given g^ℓ and X_1^n . For each ℓ , let us choose $g^{\ell^*(k)}$ as follows: $\ell^*(k) \in \arg \min_{\ell=1, \dots, N} \|g^\ell - f^k\|_2$. We then have the upper bound

$$I_{X_1^n}(\Theta; Y_1^n) \leq \frac{1}{M} \sum_{k=1}^M \left\{ \log N + \frac{n}{2} \|g^{\ell^*(k)} - f^k\|_n^2 \right\}.$$

Taking expectations over X_1^n , we obtain

$$\begin{aligned} \mathbb{E}_{X_1^n} [I_{X_1^n}(\Theta; Y_1^n)] &\leq \frac{1}{M} \sum_{k=1}^M \left\{ \log N + \frac{n}{2} \mathbb{E}_{X_1^n} [\|g^{\ell^*(k)} - f^k\|_n^2] \right\} \\ &\leq \log N + \frac{n}{2} \varepsilon^2, \end{aligned}$$

where the final inequality follows from the choice of our covering set.

From this point, we can follow the same steps as Yang and Barron (1999). The polynomial scaling (13) of the metric entropy guarantees that their conditions are satisfied, and we conclude that the minimax error is lower bounded by any $\delta_n > 0$ such that $n\delta_n^2 \geq C \log N(\delta_n; \mathcal{F}, \|\cdot\|_2)$. From Lemma 4 and the assumed scaling (13), it is equivalent to solve the equation

$$n\delta_n^2 \geq C \left\{ s \log(d/s) + s(\sqrt{s}/\delta_n)^{1/\alpha} \right\},$$

from which some algebra yields $\delta_n^2 = C \left\{ \frac{s \log(d/s)}{n} + s \left(\frac{1}{n}\right)^{\frac{2\alpha}{2\alpha+1}} \right\}$ as a suitable choice.

4.3 Proof of Theorem 3

Recall the definition of $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$ and $\mathcal{H}(S, B)$ from Section 3.5; note that it guarantees that $\|f^*\|_\infty \leq B$. In order to establish upper bounds on the minimax rate in $L^2(\mathbb{P})$ -error over $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$, we analyze a least-squares estimator—albeit *not* the same as the original M-estimator (6)—constrained to $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$, namely

$$\hat{f} \in \arg \min_{f \in \mathcal{F}_{d,s,\mathcal{H}}^*(B)} \sum_{i=1}^n (y_i - \bar{y}_n - f(x_i))^2. \quad (30)$$

Since our goal is to upper bound the minimax rate in $L^2(\mathbb{P})$ error, it is sufficient to upper bound the $L^2(\mathbb{P})$ -norm of $\hat{f} - f^*$ where \hat{f} is any solution to (30). The proof shares many steps with the proof of Theorem 1. First, the same reasoning shows that the error $\hat{\Delta} := \hat{f} - f^*$ satisfies the basic inequality

$$\frac{1}{n} \sum_{i=1}^n \hat{\Delta}^2(x_i) \leq \frac{2}{n} \left| \sum_{i=1}^n w_i \hat{\Delta}(x_i) \right| + |\bar{y}_n - \bar{f}| \left| \frac{1}{n} \sum_{i=1}^n \hat{\Delta}(x_i) \right|.$$

Recall the definition (14) of the critical rate δ_n . Once again, we first control the term error due to estimating the mean $|\bar{y}_n - \bar{f}| = \left| \frac{1}{n} \sum_{i=1}^n (y_i - \bar{f}) \right|$. Since $|f^*(x_i)|$ is at most B and w_i is standard Gaussian and independent, the random variable $y_i - \bar{f} = f^*(x_i) + w_i$ is sub-Gaussian with parameter $\sqrt{B^2 + 1}$. The samples are all i.i.d., so that by standard sub-Gaussian tail bounds, we have

$$\mathbb{P}[|\bar{y}_n - \bar{f}| > t] \leq 2 \exp\left(-\frac{nt^2}{2(B^2 + 1)}\right).$$

Setting $\mathcal{A}(\delta_n) = \{|\bar{y}_n - \bar{f}| \leq B\delta_n\}$, it is clear that

$$\mathbb{P}[\mathcal{A}(\delta_n)] \geq 1 - 2 \exp\left(-\frac{n\delta_n^2}{4}\right).$$

For the remainder of the proof, we condition on the event $\mathcal{A}(\delta_n)$, in which case Equation (17) simplifies to

$$\frac{1}{2} \|\hat{\Delta}\|_n^2 \leq \left| \frac{1}{n} \sum_{i=1}^n w_i \hat{\Delta}(x_i) \right| + B\delta_n \|\hat{\Delta}\|_n. \quad (31)$$

Here we have used the fact that $\left| \frac{1}{n} \sum_{i=1}^n \hat{\Delta}(x_i) \right| \leq \|\hat{\Delta}\|_n$, by the Cauchy-Schwartz inequality.

Now we control the Gaussian complexity term $\left| \frac{1}{n} \sum_{i=1}^n w_i \hat{\Delta}(x_i) \right|$. For any fixed subset S , define the random variable

$$\hat{Z}_n(w, t; \mathcal{H}(S, 2B)) := \sup_{\substack{\Delta \in \mathcal{H}(S, 2B) \\ \|\Delta\|_n \leq t}} \left| \frac{1}{n} \sum_{i=1}^n w_i \Delta(x_i) \right|. \quad (32)$$

We first bound this random variable for a fixed subset S of size $2s$, and then take the union bound over all $\binom{d}{2s}$ possible subsets.

Lemma 5 *Assume that the RKHS \mathcal{H} has eigenvalues $(\mu_k)_{k=1}^\infty$ that satisfy $\mu_k \simeq k^{-2\alpha}$ and eigenfunctions such that $\|\phi_k\|_\infty \leq C$. Then we have*

$$\mathbb{P}[\exists t > 0 \text{ such that } \widehat{Z}_n(w, t; \mathcal{H}(S, 2B)) \geq 16BC \sqrt{\frac{s^{1/\alpha} \log s}{n}} + 3t\delta_n] \leq c_1 \exp(-9n\delta_n^2).$$

The proof of Lemma 5 is provided Appendix F.1. Returning to inequality (31), we note that by definition,

$$\frac{2}{n} \left| \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| \leq \max_{|S|=2s} \widehat{Z}_n(w, \|\widehat{\Delta}\|_n; \mathcal{H}(S, 2B)).$$

Lemma 5 combined with the union bound implies that

$$\max_{|S|=2s} \widehat{Z}_n(w, \|\widehat{\Delta}\|_n; \mathcal{H}(S, 2B)) \leq 16BC \sqrt{\frac{s^{1/\alpha} \log s}{n}} + 3\delta_n \|\widehat{\Delta}\|_n$$

with probability at least $1 - c_1 \binom{d}{2s} \exp(-3n\delta_n^2)$. Our choice (14) of δ_n ensures that this probability is at least $1 - c_1 \exp(-c_2 n \delta_n^2)$. Combined with the basic inequality (31), we conclude that

$$\|\widehat{\Delta}\|_n^2 \leq 32BC \sqrt{\frac{s^{1/\alpha} \log s}{n}} + 7B\delta_n \|\widehat{\Delta}\|_n \quad (33)$$

with probability $1 - c_1 \exp(-c_2 n \delta_n^2)$.

By definition (14) of δ_n , the bound (33) implies that $\|\widehat{\Delta}\|_n = O(\delta_n)$ with high probability. In order to translate this claim into a bound on $\|\widehat{\Delta}\|_2$, we require the following result:

Lemma 6 *There exist universal constants (c, c_1, c_2) such that for all $t \geq c\delta_n$, we have*

$$\frac{\|g\|_2}{2} \leq \|g\|_n \leq \frac{3}{2} \|g\|_2 \quad \text{for all } g \in \mathcal{H}(S, 2B) \text{ with } \|g\|_2 \geq t \quad (34)$$

with probability at least $1 - c_1 \exp(-c_2 n t^2)$.

Proof The bound (34) follows by applying Lemma 7 in Appendix A with $\mathcal{G} = \mathcal{H}(S, 2B)$ and $b = 2B$. The critical radius from equation (35) needs to satisfy the relation $Q_{w,n}(\varepsilon_n; \mathcal{H}(S, 2B)) \leq \frac{\varepsilon_n^2}{40}$. From Lemma 11, the choice $\varepsilon_n^2 = 320BC \sqrt{\frac{s^{1/\alpha} \log s}{n}}$ satisfies this relation. By definition (14) of δ_n , we have $\delta_n \geq c\varepsilon_n$ for some universal constant c , which completes the proof. \blacksquare

This lemma implies that with probability at least $1 - c_1 \exp(-c_2 B n \delta_n^2)$, we have $\|\widehat{\Delta}\|_2 \leq 2\|\widehat{\Delta}\|_n + C\delta_n$. Combined with our earlier upper bound on $\|\widehat{\Delta}\|_n$, this completes the proof of Theorem 3.

5. Discussion

In this paper, we have studied estimation in the class of sparse additive models in which each univariate function lies within a reproducing kernel Hilbert space. In conjunction, Theorems 1 and 2

provide a precise characterization of the minimax-optimal rates for estimating f^* in the $L^2(\mathbb{P})$ -norm for various kernel classes with bounded univariate functions. These classes include finite-rank kernels (with logarithmic metric entropy), as well as kernels with polynomially decaying eigenvalues (and hence polynomial metric entropy). In order to establish achievable rates, we analyzed a simple M -estimator based on regularizing the least-squares loss with two kinds of ℓ_1 -based norms, one defined by the univariate Hilbert norm and the other by the univariate empirical norm. On the other hand, we obtained our lower bounds by a combination of approximation-theoretic and information-theoretic techniques.

An important feature of our analysis is we assume only that each univariate function is bounded, but do not assume that the multivariate function class is bounded. As discussed in Section 3.5, imposing a global boundedness condition in the high-dimensional setting can lead to a substantially smaller function classes; for instance, for Sobolev classes and sparsity $s = \Omega(\sqrt{n})$, Theorem 3 shows that it is possible to obtain much faster rates than the optimal rates for the class of sparse additive models with univariate functions bounded. Theorem 3 in our paper shows that the rates obtained under global boundedness conditions are not minimax optimal for Sobolev spaces in the regime $s = \Omega(\sqrt{n})$.

There are a number of ways in which this work could be extended. Our work considered only a hard sparsity model, in which at most s co-ordinate functions were non-zero, whereas it could be realistic to use a “soft” sparsity model involving ℓ_q -norms. Some recent work by Suzuki and Sugiyama (2012) has studied some extensions of this type. In addition, the analysis here was based on assuming independence of the covariates x_j , $j = 1, 2, \dots, d$; it would be interesting to investigate the case when the random variables are endowed with some correlation structure. One might expect some changes in the optimal rates, particularly if many of the variables are strongly dependent. Finally, this work considered only the function class consisting of sums of co-ordinate functions, whereas a natural extension would be to consider nested non-parametric classes formed of sums over hierarchies of subsets of variables.

Acknowledgments

This work was partially supported by NSF grants DMS-0605165 and DMS-0907632 to MJW and BY. In addition, BY was partially supported by the NSF grant SES-0835531 (CDI), NSF grant 0939370 (CCF), and the ARO grant (W911NF-11-1-0114). MJW was also partially supported AFOSR Grant FA9550-09-1-0466. During this work, GR was financially supported by a Berkeley Graduate Fellowship. The authors would like to thank Arash Amini and Sara van de Geer for helpful discussions.

Appendix A. A General Result On Equivalence Of $L^2(\mathbb{P})$ And $L^2(\mathbb{P}_n)$ Norms

Since it is required in a number of our proofs, we begin by stating and proving a general result that provides uniform control on the difference between the empirical $\|\cdot\|_n$ and population $\|\cdot\|_2$ norms over a uniformly bounded function class \mathcal{G} . We impose two conditions on this class:

- (a) it is uniformly bounded, meaning that there is some $b \geq 1$ such that $\|g\|_\infty \leq b$ for all $g \in \mathcal{G}$.
- (b) it is star-shaped, meaning that if $g \in \mathcal{G}$, then $\lambda g \in \mathcal{G}$ for all $\lambda \in [0, 1]$.

For each co-ordinate, the Hilbert ball $\mathbb{B}_{\mathcal{H}}(2)$ satisfies both of these conditions; we use $\mathcal{G} = \mathbb{B}_{\mathcal{H}}(2)$. (To be clear, we cannot apply this result to the multivariate function class $\mathcal{F}_{d,s,\mathcal{H}}$, since it is not uniformly bounded.)

Let $\{\sigma_i\}_{i=1}^n$ be an i.i.d. sequence of Rademacher variables, and let $\{x_i\}_{i=1}^n$ be an i.i.d. sequence of variables from \mathcal{X} , drawn according to some distribution \mathbb{Q} . For each $t > 0$, we define the local Rademacher complexity

$$Q_{\sigma,n}(t, \mathcal{G}) := \mathbb{E}_{x,\sigma} \left[\sup_{\substack{\|g\|_2 \leq t \\ g \in \mathcal{G}}} \frac{1}{n} \sum_{i=1}^n \sigma_i g(x_i) \right]$$

We let ε_n denote the smallest solution (of size at least $1/\sqrt{n}$) to the inequality

$$Q_{\sigma,n}(\varepsilon_n, \mathcal{G}) = \frac{\varepsilon_n^2}{40}, \quad (35)$$

where our scaling by the constant 40 is for later theoretical convenience. Such an ε_n exists, because the star-shaped property implies that the function $Q_{\sigma,n}(t, \mathcal{G})/t$ is non-increasing in t . This quantity corresponds to the critical rate associated with the population Rademacher complexity. For any $t \geq \varepsilon_n$, we define the event $\mathcal{E}(t) := \left\{ \sup_{\substack{g \in \mathcal{G} \\ \|g\|_2 \leq t}} \left| \|g\|_n - \|g\|_2 \right| \geq \frac{bt}{2} \right\}$.

Lemma 7 *Suppose that $\|g\|_\infty \leq b$ for all $g \in \mathcal{G}$. Then there exist universal constants (c_1, c_2) such that for any $t \geq \varepsilon_n$,*

$$\mathbb{P}[\mathcal{E}(t)] \leq c_1 \exp(-c_2 n t^2).$$

In addition, for any $g \in \mathcal{G}$ with $\|g\|_2 \geq t$, we have $\|g\|_n \leq \|g\|_2(1 + \frac{b}{2})$, and moreover, for all $g \in \mathcal{G}$ with $\|g\|_2 \geq bt$, we have

$$\frac{1}{2} \|g\|_2 \leq \|g\|_n \leq \frac{3}{2} \|g\|_2, \quad (36)$$

both with probability at least $1 - c_1 \exp(-c_2 n t^2)$.

Lemma 7 follows from a relatively straightforward adaptation of known results (e.g., Lemma 5.16 in van de Geer, 2000 and Theorem 2.1 in Bartlett et al., 2005), so we omit the proof details here.

Appendix B. Proof of Lemma 1

The proof of this lemma is based on peeling and weighting techniques from empirical process theory (Alexander, 1987; van de Geer, 2000) combined with results on the local Rademacher and Gaussian complexities of kernel classes (Bartlett et al., 2005; Mendelson, 2002). For each univariate Hilbert space $\mathcal{H}_j = \mathcal{H}$, let us introduce the random variables

$$\widehat{Z}_n(w; t; \mathcal{H}) := \sup_{\substack{\|g_j\|_{\mathcal{H}} \leq 1 \\ \|g_j\|_n \leq t}} \left| \frac{1}{n} \sum_{i=1}^n w_i g_j(x_{ij}) \right|, \quad \text{and} \quad Z_n(w; t; \mathcal{H}) := \mathbb{E}_x \left[\sup_{\substack{\|g_j\|_{\mathcal{H}} \leq 1 \\ \|g_j\|_2 \leq t}} \left| \frac{1}{n} \sum_{i=1}^n w_i g_j(x_{ij}) \right| \right], \quad (37)$$

where $w_i \sim N(0, 1)$ are i.i.d. standard normal. The empirical and population Gaussian complexities are given by

$$\widehat{Q}_{w,n}(t, \mathcal{H}) := \mathbb{E}_w [\widehat{Z}_n(w; t; \mathcal{H})] \quad \text{and} \quad Q_{w,n}(t, \mathcal{H}) := \mathbb{E}_w [Z_n(w; t; \mathcal{H})].$$

For future reference, we note that in the case of a univariate Hilbert space \mathcal{H} with eigenvalues $\{\mu_k\}_{k=1}^\infty$, results in Mendelson (2002) imply that there are universal constants $c_\ell \leq c_u$ such that for all $t^2 \geq 1/n$, we have

$$\frac{c_\ell}{\sqrt{n}} \left[\sum_{k=1}^\infty \min\{t^2, \mu_k\} \right]^{1/2} \leq Q_{w,n}(t, \mathcal{H}) \leq \frac{c_u}{\sqrt{n}} \left[\sum_{k=1}^\infty \min\{t^2, \mu_k\} \right]^{1/2}, \quad (38)$$

for all j . The same bounds hold for the local Rademacher complexity in our special case of reproducing kernel Hilbert spaces.

Let $\widehat{v}_{n,j} > 0$ denote the smallest positive solution r of the inequality

$$\widehat{Q}_{w,n}(r, \mathcal{H}) \leq 4r^2. \quad (39)$$

The function $\widehat{Q}_{w,n}(r, \mathcal{H})$ defines the local Gaussian complexity of the kernel class in co-ordinate j . Recall the bounds (38) that apply to both the empirical and population Gaussian complexities. Recall that the critical univariate rate v_n is defined in terms of the population Gaussian complexity (see Equation (8)).

B.1 Some Auxiliary Results

In order to prove Lemma 1, we also need some auxiliary results, stated below as Lemmas 8 and 9.

Lemma 8 *For any function class \mathcal{G} and all $\delta \geq 0$, we have*

$$\mathbb{P} \left[|\widehat{Z}_n(w, t, \mathcal{G}) - \widehat{Q}_{w,n}(t, \mathcal{G})| \geq \delta t \right] \leq 2 \exp \left(-\frac{n\delta^2}{2} \right), \quad \text{and} \quad (40)$$

$$\mathbb{P} \left[|Z_n(w, t, \mathcal{G}) - Q_{w,n}(t, \mathcal{G})| \geq \delta t \right] \leq 2 \exp \left(-\frac{n\delta^2}{2} \right). \quad (41)$$

Proof We have

$$|\widehat{Z}_n(w, t, \mathcal{G}) - \widehat{Z}_n(w', t, \mathcal{G})| \leq \sup_{\substack{g \in \mathcal{G} \\ \|g\|_n \leq t}} \frac{1}{n} \left| \sum_{i=1}^n (w_i - w'_i) g(x_i) \right| \leq \frac{t}{\sqrt{n}} \|w - w'\|_2,$$

showing that $\widehat{Z}_n(w, t, \mathcal{G})$ is $\frac{t}{\sqrt{n}}$ -Lipschitz with respect to the ℓ_2 norm. Consequently, concentration for Lipschitz functions of Gaussian random variables (see Ledoux, 2001) yields the tail bound (40). Turning to the quantity $Z_n(w, t, \mathcal{H})$, a similar argument yields that

$$\begin{aligned} |Z_n(w, t, \mathcal{G}) - Z_n(w', t, \mathcal{G})| &\leq \mathbb{E}_x \left[\sup_{\substack{g \in \mathcal{G} \\ \|g\|_2 \leq t}} \frac{1}{n} \left| \sum_{i=1}^n (w_i - w'_i) g(x_i) \right| \right] \\ &\leq \sup_{\substack{g \in \mathcal{G} \\ \|g\|_2 \leq t}} \mathbb{E}_x \left[\left(\frac{1}{n} \sum_{i=1}^n g^2(x_i) \right)^{1/2} \right] \|w - w'\|_2 \leq \frac{t}{\sqrt{n}} \|w - w'\|_2, \end{aligned}$$

where the final step uses Jensen's inequality and the fact that $\mathbb{E}_x[g^2(x_i)] \leq t^2$ for all $i = 1, \dots, n$. The same reasoning then yields the tail bound (41). \blacksquare

Our second lemma involves the event $\mathcal{D}(\gamma_n) := \{\widehat{v}_{n,j} \leq \gamma_n, \text{ for all } j = 1, 2, \dots, d\}$, where we recall the definition (39) of $\widehat{v}_{n,j}$, and that $\gamma_n := \kappa \max\{v_n, \sqrt{\frac{\log d}{n}}\}$.

Lemma 9 *For all $1 \leq j \leq d$, we have*

$$\mathbb{P}[\widehat{v}_{n,j} \leq \gamma_n] \geq 1 - c_1 \exp(-c_2 n \gamma_n^2).$$

Proof

We first bound the probability of the event $\{\widehat{v}_{n,j} > \gamma_n\}$ for a fixed \mathcal{H}_j . Let $g \in \mathbb{B}_{\mathcal{H}_j}(1)$ be any function such that $\|g\|_2 > t \geq v_n$. Then conditioned on the sandwich relation (36) with $b = 1$, we are guaranteed that $\|g\|_n > \frac{t}{2}$. Taking the contrapositive, we conclude that $\|g\|_n \leq \frac{t}{2}$ implies $\|g\|_2 \leq t$, and hence that $\widehat{Z}_n(w, t/2, \mathcal{H}) \leq Z_n(w, t, \mathcal{H})$ for all $t \geq v_n$, under the stated conditioning.

For any $t \geq v_n$, the inequalities (36), (40) and (41) hold with probability at least $1 - c_1 \exp(-c_2 n t^2)$. Conditioning on these inequalities, we can set $t = \gamma_n > v_n$, and thereby obtain

$$\begin{aligned} \widehat{Q}_{w,n}(\gamma_n, \mathcal{H}) &\stackrel{(a)}{\leq} \widehat{Z}_n(w, \gamma_n, \mathcal{H}) + \gamma_n^2 \\ &\stackrel{(b)}{\leq} Z_n(w, 2\gamma_n, \mathcal{H}) + \gamma_n^2 \\ &\stackrel{(c)}{\leq} Q_{w,n}(2\gamma_n, \mathcal{H}) + 2\gamma_n^2 \\ &\stackrel{(d)}{\leq} 4\gamma_n^2, \end{aligned}$$

where inequality (a) follows from the bound (40), inequality (b) follows the initial argument, inequality (c) follows from the bound (41), and inequality (d) follows since $2\gamma_n > \varepsilon_n$ and the definition of ε_n .

By the definition of $\widehat{v}_{n,j}$ as the minimal t such that $\widehat{Q}_{w,n}(t, \mathcal{H}) \leq 4t^2$, we conclude that for each fixed $j = 1, \dots, n$, we have $\widehat{v}_{n,j} \leq \gamma_n$ with probability at least $1 - c_1 \exp(-c_2 n \gamma_n^2)$. Finally, the uniformity over $j = 1, 2, \dots, d$ follows from the union bound and our choice of $\gamma_n \geq \kappa \sqrt{\frac{\log d}{n}}$. \blacksquare

B.2 Main Argument To Prove Lemma 1

We can now proceed with the proof of Lemma 1. Combining Lemma 9 with the union bound over $j = 1, 2, \dots, d$, we conclude that that

$$\mathbb{P}[\mathcal{D}(\gamma_n)] \geq 1 - c_1 \exp(-c_2 n \gamma_n^2),$$

as long as $c_2 \geq 1$. For the remainder of our proofs, we condition on the event $\mathcal{D}(\gamma_n)$. In particular, our goal is to prove that

$$\left| \frac{1}{n} \sum_{i=1}^n w_i f_j(x_{ij}) \right| \leq C \{ \gamma_n^2 \|f_j\|_{\mathcal{H}} + \gamma_n \|f_j\|_n \} \quad \text{for all } f_j \in \mathcal{H} \quad (42)$$

with probability greater than $1 - c_1 \exp(-c_2 n \gamma_n^2)$. By combining this result with our choice of γ_n and the union bound, the claimed bound then follows on $\mathbb{P}[\mathcal{T}(\gamma_n)]$.

If $f_j = 0$, then the claim (42) is trivial. Otherwise we renormalize f_j by defining $g_j := f_j / \|f_j\|_{\mathcal{H}}$, and we write

$$\frac{1}{n} \sum_{i=1}^n w_i f_j(x_{ij}) = \|f_j\|_{\mathcal{H}} \frac{1}{n} \sum_{i=1}^n w_i g_j(x_{ij}) \leq \|f_j\|_{\mathcal{H}} \widehat{\mathcal{Z}}_n(w; \|g_j\|_n, \mathcal{H}),$$

where the final inequality uses the definition (37), and the fact that $\|g_j\|_{\mathcal{H}} = 1$. We now split the analysis into two cases: (1) $\|g_j\|_n \leq \gamma_n$, and (2) $\|g_j\|_n > \gamma_n$.

Case 1: $\|g_j\|_n \leq \gamma_n$. In this case, it suffices to upper bound the quantity $\widehat{\mathcal{Z}}_n(w; \gamma_n, \mathcal{H})$. Note that $\|g_j\|_{\mathcal{H}} = 1$ and recall definition (37) of the random variable $\widehat{\mathcal{Z}}_n$. On one hand, since $\gamma_n \geq \widehat{v}_{n,j}$ by Lemma 9, the definition of $\widehat{v}_{n,j}$ implies that $\widehat{\mathcal{Q}}_{w,n}(\gamma_n, \mathcal{H}) \leq 4\gamma_n^2$, and hence

$$\mathbb{E}[\widehat{\mathcal{Z}}_n(w; \gamma_n; \mathcal{H})] = \widehat{\mathcal{Q}}_{w,n}(\gamma_n; \mathcal{H}) \leq 4\gamma_n^2.$$

Applying the bound (40) from Lemma 8 with $\delta = \gamma_n = t$, we conclude that $\widehat{\mathcal{Z}}_n(w; \gamma_n; \mathcal{H}) \leq C\gamma_n^2$ with probability at least $1 - c_1 \exp\{-c_2 n \gamma_n^2\}$, which completes the proof in the case where $\|g_j\|_n \leq \gamma_n$.

Case 2: $\|g_j\|_n > \gamma_n$. In this case, we study the random variable $\widehat{\mathcal{Z}}_n(w; r_j; \mathcal{H})$ for some $r_j > \gamma_n$. Our intermediate goal is to prove the bound

$$\mathbb{P}\left[\widehat{\mathcal{Z}}_n(w; r_j; \mathcal{H}) \geq C r_j \gamma_n\right] \leq c_1 \exp\{-c_2 n \gamma_n^2\}. \quad (43)$$

Applying the bound (40) with $t = r_j$ and $\delta = \gamma_n$, we are guaranteed an upper bound of the form $\widehat{\mathcal{Z}}_n(w; r_j; \mathcal{H}) \leq \widehat{\mathcal{Q}}_{w,n}(r_j, \mathcal{H}) + r_j \gamma_n$ with probability at least $1 - c_1 \exp(-c_2 n \gamma_n^2)$. In order to complete the proof, we need to show that $\widehat{\mathcal{Q}}_{w,n}(r_j, \mathcal{H}) \leq r_j \gamma_n$. Since $r_j > \gamma_n > \widehat{v}_{n,j}$, we have

$$\widehat{\mathcal{Q}}_{w,n}(r_j, \mathcal{H}) = \frac{r_j}{\widehat{v}_{n,j}} \mathbb{E}_w \left[\sup_{\substack{\|g_j\|_n \leq \widehat{v}_{n,j} \\ \|g_j\|_{\mathcal{H}} \leq \frac{\widehat{v}_{n,j}}{r_j}}} \left| \frac{1}{n} \sum_{i=1}^n w_i g_j(x_{ij}) \right| \right] \leq \frac{r_j}{\widehat{v}_{n,j}} \widehat{\mathcal{Q}}_{w,n}(\widehat{v}_{n,j}, \mathcal{H}) \leq 4 r_j \widehat{v}_{n,j},$$

where the final inequality uses the fact that $\widehat{\mathcal{Q}}_{w,n}(\widehat{v}_{n,j}, \mathcal{H}) \leq 4\widehat{v}_{n,j}^2$. On the event $\mathcal{D}(\gamma_n)$ from Lemma 9, we have $\widehat{v}_{n,j} \leq \gamma_n$, from which the claim (43) follows.

We now use the bound (43) to prove the bound (42), in particular via a ‘‘peeling’’ operation over all choices of $r_j = \|f_j\|_n / \|f_j\|_{\mathcal{H}}$. (See van de Geer, 2000 for more details on these peeling arguments.) We claim that it suffices to consider $r_j \leq 1$. It is equivalent to show that $\|g_j\|_n \leq 1$ for any $g_j \in \mathbb{B}_{\mathcal{H}}(1)$. Since $\|g_j\|_{\infty} \leq \|g_j\|_{\mathcal{H}} \leq 1$, we have $\|g_j\|_n^2 = \frac{1}{n} \sum_{i=1}^n g_j^2(x_{ij}) \leq 1$, as required. Now define the event

$$\mathcal{T}_j(\gamma_n) := \left\{ \exists f_j \in \mathbb{B}_{\mathcal{H}}(1) \mid \left| \frac{1}{n} \sum_{i=1}^n w_i f_j(x_{ij}) \right| > 8 \|f_j\|_{\mathcal{H}} \gamma_n \frac{\|f_j\|_n}{\|f_j\|_{\mathcal{H}}}, \text{ and } \frac{\|f_j\|_n}{\|f_j\|_{\mathcal{H}}} \in (\gamma_n, 1] \right\}.$$

and the sets $S_m := \{2^{m-1} \gamma_n \leq \frac{\|f_j\|_n}{\|f_j\|_{\mathcal{H}}} \leq 2^m \gamma_n\}$ for $m = 1, 2, \dots, M$. By choosing $M = 2 \log_2(1/\gamma_n)$, we ensure that $2^M \gamma_n \geq 1$, and hence that if the event $\mathcal{T}_j(\gamma_n)$ occurs, then it must occur for function

f_j belonging to some S_m , so that we have a function f_j such that $\frac{\|f_j\|_n}{\|f_j\|_{\mathcal{H}}} \leq t_m := 2^m \gamma_n$, and

$$\left| \frac{1}{n} \sum_{i=1}^n w_i f_j(x_{ij}) \right| > 8 \|f_j\|_{\mathcal{H}} \gamma_n \frac{\|f_j\|_n}{\|f_j\|_{\mathcal{H}}} \geq C \|f_j\|_{\mathcal{H}} t_m,$$

which implies that $\widehat{Z}_n(w; t_m, \mathcal{H}) \geq 4t_m$. Consequently, by union bound and the tail bound (43), we have

$$\mathbb{P}[\mathcal{T}_j(\gamma_n)] \leq M c_1 \exp\{-c_2 n \gamma_n^2\} \leq c_1 \exp\{-c'_2 n \gamma_n^2\}$$

by the condition $n \gamma_n^2 = \Omega(\log(1/\gamma_n))$, which completes the proof.

Appendix C. Proof of Lemma 2

Define the function

$$\widetilde{\mathcal{L}}(\Delta) := \frac{1}{2n} \sum_{i=1}^n (w_i + \bar{f} + \bar{y}_n - \Delta(x_i))^2 + \lambda_n \|f^* + \Delta\|_{n,1} + \rho_n \|f^* + \Delta\|_{\mathcal{H},1}$$

and note that by definition of our M -estimator, the error function $\widehat{\Delta} := \widehat{f} - f^*$ minimizes $\widetilde{\mathcal{L}}$. From the inequality $\widetilde{\mathcal{L}}(\widehat{\Delta}) \leq \widetilde{\mathcal{L}}(0)$, we obtain the upper bound $\frac{1}{2} \|\widehat{\Delta}\|_n^2 \leq T_1 + T_2$, where

$$\begin{aligned} T_1 &:= \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| + |\bar{y}_n - \bar{f}| \left| \frac{1}{n} \sum_{i=1}^n \widehat{\Delta}(x_i) \right|, \quad \text{and} \\ T_2 &:= \lambda_n \sum_{j=1}^d \{ \|f_j^*\|_n - \|f_j^* + \widehat{\Delta}_j\|_n \} + \rho_n \sum_{j=1}^d \{ \|f_j^*\|_{\mathcal{H}} - \|f_j^* + \widehat{\Delta}_j\|_{\mathcal{H}} \}. \end{aligned}$$

Conditioned on the event $\mathcal{C}(\gamma_n)$, we have the bound $|\bar{y}_n - \bar{f}| \left| \frac{1}{n} \sum_{i=1}^n \widehat{\Delta}(x_i) \right| \leq \sqrt{s} \gamma_n \|\widehat{\Delta}\|_n$, and hence $\frac{1}{2} \|\widehat{\Delta}\|_n^2 \leq T_2 + \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| + \sqrt{s} \gamma_n \|\widehat{\Delta}\|_n$, or equivalently

$$0 \leq \frac{1}{2} (\|\widehat{\Delta}\|_n - \sqrt{s} \gamma_n)^2 \leq T_2 + \left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| + \frac{1}{2} s \gamma_n^2. \quad (44)$$

It remains to control the term T_2 . On one hand, for any $j \in S^c$, we have

$$\|f_j^*\|_n - \|f_j^* + \widehat{\Delta}_j\|_n = -\|\widehat{\Delta}_j\|_n, \quad \text{and} \quad \|f_j^*\|_{\mathcal{H}} - \|f_j^* + \widehat{\Delta}_j\|_{\mathcal{H}} = -\|\widehat{\Delta}_j\|_{\mathcal{H}}.$$

On the other hand, for any $j \in S$, the triangle inequality yields $\|f_j^*\|_n - \|f_j^* + \widehat{\Delta}_j\|_n \leq \|\widehat{\Delta}_j\|_n$, with a similar inequality for the terms involving $\|\cdot\|_{\mathcal{H}}$. Combined with the bound (44), we conclude that

$$0 \leq \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) + \lambda_n \{ \|\widehat{\Delta}_S\|_{n,1} - \|\widehat{\Delta}_{S^c}\|_{n,1} \} + \rho_n \{ \|\widehat{\Delta}_S\|_{\mathcal{H},1} - \|\widehat{\Delta}_{S^c}\|_{\mathcal{H},1} \} + \frac{1}{2} s \gamma_n^2. \quad (45)$$

Recalling our conditioning on the event $\mathcal{T}(\gamma_n)$, by Lemma 1, we have the upper bound

$$\left| \frac{1}{n} \sum_{i=1}^n w_i \widehat{\Delta}(x_i) \right| \leq 8 \{ \gamma_n \|\widehat{\Delta}\|_{n,1} + \gamma_n^2 \|\widehat{\Delta}\|_{\mathcal{H},1} \}.$$

Combining with the inequality (45) yields

$$\begin{aligned} 0 &\leq 8 \{ \gamma_n \|\widehat{\Delta}\|_{n,1} + \gamma_n^2 \|\widehat{\Delta}\|_{\mathcal{H},1} \} + \lambda_n \{ \|\widehat{\Delta}_S\|_{n,1} - \|\widehat{\Delta}_{S^c}\|_{n,1} \} + \rho_n \{ \|\widehat{\Delta}_S\|_{\mathcal{H},1} - \|\widehat{\Delta}_{S^c}\|_{\mathcal{H},1} \} + \frac{1}{2} s \gamma_n^2 \\ &\leq \frac{\lambda_n}{2} \|\widehat{\Delta}\|_{n,1} + \frac{\rho_n}{2} \|\widehat{\Delta}\|_{\mathcal{H},1} + \lambda_n \{ \|\widehat{\Delta}_S\|_{n,1} - \|\widehat{\Delta}_{S^c}\|_{n,1} \} + \rho_n \{ \|\widehat{\Delta}_S\|_{\mathcal{H},1} - \|\widehat{\Delta}_{S^c}\|_{\mathcal{H},1} \} + \frac{1}{2} s \gamma_n^2, \end{aligned}$$

where we have recalled our choices of (λ_n, ρ_n) . Finally, re-arranging terms yields the claim (19).

Appendix D. Proof of Lemma 3

Recalling the definitions (26), (27) and (28) of the function class $\mathcal{G}(\lambda_n, \rho_n)$ and the critical radius $\tilde{\delta}_n$ from Equation (24), we define the function class $\mathcal{G}'(\lambda_n, \rho_n, \tilde{\delta}_n) := \{h \in \mathcal{G}(\lambda_n, \rho_n) \mid \|h\|_2 = \tilde{\delta}_n\}$, and the alternative event

$$\mathcal{B}'(\lambda_n, \rho_n) := \{ \|h\|_n^2 \geq \tilde{\delta}_n^2/2 \quad \text{for all } h \in \mathcal{G}'(\lambda_n, \rho_n, \tilde{\delta}_n) \}.$$

We claim that it suffices to show that $\mathcal{B}'(\lambda_n, \rho_n)$ holds with probability at least $1 - c_1 \exp(-c_2 n \gamma_n^2)$. Indeed, given an arbitrary non-zero function $g \in \mathcal{G}(\lambda_n, \rho_n)$, consider the rescaled function $h = \frac{\tilde{\delta}_n}{\|g\|_2} g$. Since $g \in \mathcal{G}(\lambda_n, \rho_n)$ and $\mathcal{G}(\lambda_n, \rho_n)$ is star-shaped, we have $h \in \mathcal{G}(\lambda_n, \rho_n)$, and also $\|h\|_2 = \tilde{\delta}_n$ by construction. Consequently, when the event $\mathcal{B}'(\lambda_n, \rho_n)$ holds, we have $\|h\|_n^2 \geq \tilde{\delta}_n^2/2$, or equivalently $\|g\|_n^2 \geq \|g\|_2^2/2$, showing that $\mathcal{B}(\lambda_n, \rho_n)$ holds. Accordingly, the remainder of the proof is devoted to showing that $\mathcal{B}'(\lambda_n, \rho_n)$ holds with probability greater than $1 - c_1 \exp(-c_2 n \gamma_n^2)$. Alternatively, if we define the random variable $Z_n(\mathcal{G}') := \sup_{f \in \mathcal{G}'} \{ \tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n f^2(x_i) \}$, then it suffices to show that $Z_n(\mathcal{G}') \leq \tilde{\delta}_n^2/2$ with high probability.

Recall from Section 4.2.4 the definition of a covering set; here we use the notion of a proper covering, which restricts the covering to use only members of the set \mathcal{G} . Letting $N_{\text{pr}}(\varepsilon; \mathcal{G}, \rho)$ denote the proper covering number, it can be shown that $N_{\text{pr}}(\varepsilon; \mathcal{G}, \rho) \leq N(\varepsilon; \mathcal{G}, \rho) \leq N_{\text{pr}}(\varepsilon/2; \mathcal{G}, \rho)$. Now let g^1, \dots, g^N be a minimal $\tilde{\delta}_n/8$ -proper covering of \mathcal{G}' in the $L^2(\mathbb{P}_n)$ -norm, so that for all $f \in \mathcal{G}'$, there exists $g = g^k \in \mathcal{G}'$ such that $\|f - g\|_n \leq \tilde{\delta}_n/8$. We can then write

$$\tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n f^2(x_i) = \left\{ \tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n g^2(x_i) \right\} + \left\{ \frac{1}{n} \sum_{i=1}^n (g^2(x_i) - f^2(x_i)) \right\}.$$

By the Cauchy-Schwartz inequality, we have

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (g^2(x_i) - f^2(x_i)) &= \frac{1}{n} \sum_{i=1}^n (g(x_i) - f(x_i))(g(x_i) + f(x_i)) \\ &\leq \sqrt{\frac{1}{n} \sum_{i=1}^n (g(x_i) - f(x_i))^2} \sqrt{\frac{1}{n} \sum_{i=1}^n (f(x_i) + g(x_i))^2} \\ &= \|g - f\|_n \sqrt{\frac{1}{n} \sum_{i=1}^n (f(x_i) + g(x_i))^2}. \end{aligned}$$

By our choice of the covering, we have $\|g - f\|_n \leq \tilde{\delta}_n/8$. On the other hand, we have

$$\sqrt{\frac{1}{n} \sum_{i=1}^n (f(x_i) + g(x_i))^2} \leq \sqrt{2\|f\|_n^2 + 2\|g\|_n^2} \leq \sqrt{4\tilde{\delta}_n^2} = 2\tilde{\delta}_n,$$

where the final inequality follows since $\|f\|_n = \|g\|_n = \tilde{\delta}_n$. Overall, we have established the upper bound $\frac{1}{n} \sum_{i=1}^n (g^2(x_i) - f^2(x_i)) \leq \frac{\tilde{\delta}_n^2}{4}$, and hence shown that

$$Z_n(\mathcal{G}') \leq \max_{g^1, g^2, \dots, g^N} \left\{ \tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n (g^k(x_i)) \right\} + \frac{\tilde{\delta}_n^2}{4},$$

where $N = N_{\text{pr}}(\tilde{\delta}_n/8, \mathcal{G}', \|\cdot\|_n)$. For any g in our covering set, since $g^2(x_i) \geq 0$, we may apply a one-sided tail bound (e.g., Theorem 3.5 from Chung and Lu, 2006, or Lemma 2.1 in Einmahl and Mason, 1996) with $t = \tilde{\delta}_n^2/4$ to obtain the one-sided tail bound

$$\mathbb{P}\left[\tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n g^2(x_i) \geq \frac{\tilde{\delta}_n^2}{4}\right] \leq \exp\left(-\frac{n\tilde{\delta}_n^4}{32\mathbb{E}[g^4(x)]}\right), \quad (46)$$

where we used the upper bound $\text{var}(g^2(x)) \leq \mathbb{E}[g^4(x)]$. Next using the fact that the variables $\{g_j(x_j)\}_{j=1}^d$ are independent and zero-mean, we have

$$\begin{aligned} \mathbb{E}[g^4(x)] &= \sum_{j=1}^d \mathbb{E}[g_j^4(x_j)] + \binom{4}{2} \sum_{j \neq k} \mathbb{E}[g_j^2(x_j)] \mathbb{E}[g_k^2(x_k)] \\ &\leq 4 \sum_{j=1}^d \mathbb{E}[g_j^2(x_j)] + 6 \sum_{j=1}^d \mathbb{E}[g_j^2(x_j)] \sum_{k=1}^d \mathbb{E}[g_k^2(x_k)] \\ &\leq 4\tilde{\delta}_n^2 + 6\tilde{\delta}_n^4 \\ &\leq 10\tilde{\delta}_n^2, \end{aligned}$$

where the second inequality follows since $\|g_j\|_\infty \leq \|g_j\|_{\mathcal{H}} \leq 2$ for each j . Combining this upper bound on $\mathbb{E}[g^4(x)]$ with the earlier tail bound (46) and applying union bound yields

$$\mathbb{P}\left[\max_{k=1,2,\dots,N} \left\{ \tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n g^2(x_i) \right\} \geq \frac{\tilde{\delta}_n^2}{4}\right] \leq \exp\left(\log N_{\text{pr}}(\tilde{\delta}_n/8, \mathcal{G}', \|\cdot\|_n) - \frac{n\tilde{\delta}_n^2}{320}\right). \quad (47)$$

It remains to bound the covering entropy $\log N_{\text{pr}}(\tilde{\delta}_n/8, \mathcal{G}', \|\cdot\|_n)$. Since the proper covering entropy $\log N_{\text{pr}}(\tilde{\delta}_n/8, \mathcal{G}', \|\cdot\|_n)$ is at most $\log N(\tilde{\delta}_n/16, \mathcal{G}', \|\cdot\|_n)$, it suffices to upper bound the usual covering entropy. Viewing the samples (x_1, x_2, \dots, x_n) as fixed, let us define the zero-mean Gaussian process $\{W_g, g \in \mathcal{G}'\}$ via $W_g := \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i g(x_i)$, where the variables $\{\varepsilon_i\}_{i=1}^n$ are i.i.d. standard Gaussian variates. By construction, we have $\text{var}[(W_g - W_f)] = \|g - f\|_n^2$. Consequently, by the Sudakov minoration (see Pisier, 1989), for all $\varepsilon > 0$, we have $\varepsilon \sqrt{\log N(\varepsilon; \mathcal{G}', \|\cdot\|_n)} \leq 4\mathbb{E}_\varepsilon[\sup_{g \in \mathcal{G}'} W_g]$. Setting $\varepsilon = \tilde{\delta}_n/16$ and performing some algebra, we obtain the upper bound

$$\frac{1}{\sqrt{n}} \sqrt{\log N(\tilde{\delta}_n/16; \mathcal{G}', \|\cdot\|_n)} \leq \frac{64}{\tilde{\delta}_n} \mathbb{E}_\varepsilon\left[\sup_{g \in \mathcal{G}'} \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(x_i)\right]. \quad (48)$$

The final step is to upper bound the Gaussian complexity $\mathbb{E}_\varepsilon\left[\sup_{g \in \mathcal{G}'} \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(x_i)\right]$. In the proof of Lemma 1, we showed that for any co-ordinate $j \in \{1, 2, \dots, d\}$, the univariate Gaussian complexity is upper bounded as

$$\mathbb{E}\left[\sup_{\substack{\|g_j\|_n \leq r_j \\ \|g_j\|_{\mathcal{H}} \leq R_j}} \frac{1}{n} \sum_{i=1}^n \varepsilon_i g_j(x_{ij})\right] \leq C \{\gamma_n r_j + \gamma_n^2 R_j\}.$$

Summing across co-ordinates and recalling the fact that the constant C may change from line to line, we obtain the upper bound

$$\begin{aligned}
 \mathbb{E}_\varepsilon \left[\sup_{g \in \mathcal{G}'} \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(x_i) \right] &\leq C \sup_{g \in \mathcal{G}'} \{ \gamma_n \|g\|_{1,n} + \gamma_n^2 \|g\|_{1,\mathcal{H}} \} \\
 &\stackrel{(a)}{\leq} C \sup_{g \in \mathcal{G}'} \{ 4\gamma_n \|g_S\|_{1,n} + 4\gamma_n^2 \|g_S\|_{1,\mathcal{H}} + \frac{1}{32} s \rho_n \} \\
 &\stackrel{(b)}{\leq} C \sup_{g \in \mathcal{G}'} \{ \gamma_n \|g_S\|_{1,n} + s \rho_n \} \\
 &\stackrel{(c)}{\leq} C \sup_{g \in \mathcal{G}'} \left\{ \gamma_n [2\sqrt{s} \|g\|_2 + s \gamma_n] + s \rho_n \right\},
 \end{aligned}$$

where step (a) uses inequality (26) in the definition of \mathcal{G}' ; step (b) uses the inequality $\|g_j\|_{\mathcal{H}} \leq 2$ for each co-ordinate and hence $\|g_S\|_{1,\mathcal{H}} \leq 2s$, and our choice of regularization parameter $\rho_n \geq \gamma_n^2$; and step (c) uses inequality (27) in the definition of \mathcal{G}' . Since $\|g\|_2 = \tilde{\delta}_n$ for all $g \in \mathcal{G}'$, we have shown that

$$\mathbb{E}_\varepsilon \left[\sup_{g \in \mathcal{G}'} \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(x_i) \right] \leq C \{ s \gamma_n^2 + \sqrt{s} \gamma_n \tilde{\delta}_n + s \rho_n \} \stackrel{(d)}{\leq} C \left\{ \frac{\tilde{\delta}_n^2}{B^2} + \frac{\tilde{\delta}_n^2}{B} \right\}, \quad (49)$$

where inequality (d) follows from our choice (24) of $\tilde{\delta}_n$, and the constant B can be chosen as large as we please. In particular, by choosing B sufficiently large, and combining the bound (49) with the Sudakov bound (48), we can ensure that

$$\frac{1}{n} \log N(\tilde{\delta}_n/16; \mathcal{G}', \|\cdot\|_n) \leq \frac{\tilde{\delta}_n^2}{640}.$$

Combined with the earlier tail bound (47), we conclude that

$$\mathbb{P} \left[\max_{k=1,2,\dots,N} \left\{ \tilde{\delta}_n^2 - \frac{1}{n} \sum_{i=1}^n g^2(x_i) \right\} \geq \frac{\tilde{\delta}_n^2}{4} \right] \leq \exp \left(- \frac{n \tilde{\delta}_n^2}{640} \right),$$

which completes the proof of Lemma 3.

Appendix E. Proof of Lemma 4

In this section, we present the proofs of Lemma 4 (a) and (b).

E.1 Proof of Part (a)

Let $N = M(\frac{\delta}{\sqrt{s}}; \mathbb{B}_{\mathcal{H}}(1), \|\cdot\|_2) - 1$, and define $I = \{0, 1, \dots, N\}$. Consider the set

$$\mathfrak{S} := \{ u \in I^d \mid \|u\|_0 := \sum_{j=1}^d \mathbb{I}[u_j \neq 0] = s \}. \quad (50)$$

Note that this set has cardinality $|\mathfrak{S}| = \binom{d}{s} N^s$, since any element is defined by first choosing s co-ordinates are non-zero, and then for each co-ordinate, choosing non-zero entry from a total of N possible symbols.

For each $j = 1, \dots, d$, let $\{0, f_j^1, f_j^2, \dots, f_j^N\}$ be a δ/\sqrt{s} -packing of $\mathbb{B}_{\mathcal{H}}(1)$. Based on these packings of the univariate function classes, we can use \mathfrak{S} to index a collection of functions contained inside \mathcal{F} . In particular, any $u \in \mathfrak{S}$ uniquely defines a function $g^u = \sum_{j=1}^d g_j^{u_j} \in \mathcal{F}$, with elements

$$g_j^{u_j} = \begin{cases} f_j^{u_j} & \text{if } u_j \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Since $\|u\|_0 = s$, we are guaranteed that at most s co-ordinates of g are non-zero, so that $g \in \mathcal{F}$.

Now consider two functions g^u and h^v contained within the class $\{g^u, u \in \mathfrak{S}\}$. By definition, we have

$$\|g^u - h^v\|_2^2 = \sum_{j=1}^d \|f_j^{u_j} - f_j^{v_j}\|_2^2 \geq \frac{\delta^2}{s} \sum_{j=1}^d \mathbb{I}[u_j \neq v_j], \quad (51)$$

Consequently, it suffices to establish the existence of a ‘‘large’’ subset $\mathcal{A} \subset \mathfrak{S}$ such that the Hamming metric $\rho_H(u, v) := \sum_{j=1}^d \mathbb{I}[u_j \neq v_j]$ is at least $s/2$ for all pairs $u, v \in \mathcal{A}$, in which case we are guaranteed that $\|g - h\|_2^2 \geq \delta^2$. For any $u \in \mathfrak{S}$, we observe that

$$\left| \left\{ v \in \mathfrak{S} \mid \rho_H(u, v) \leq \frac{s}{2} \right\} \right| \leq \binom{d}{\frac{s}{2}} (N+1)^{\frac{s}{2}}.$$

This bound follows because we simply need to choose a subset of size $s/2$ where u and v agree, and the remaining $s/2$ co-ordinates can be chosen arbitrarily in $(N+1)^{\frac{s}{2}}$ ways. For a given set \mathcal{A} , we write $\rho_H(u, \mathcal{A}) \leq \frac{s}{2}$ if there exists some $v \in \mathcal{A}$ such that $\rho_H(u, v) \leq \frac{s}{2}$. Using this notation, we have

$$\left| \left\{ u \in \mathfrak{S} \mid \rho_H(u, \mathcal{A}) \leq \frac{s}{2} \right\} \right| \leq |\mathcal{A}| \binom{d}{\frac{s}{2}} (N+1)^{\frac{s}{2}} \stackrel{(a)}{<} |\mathfrak{S}|,$$

where inequality (a) follows as long as

$$|\mathcal{A}| \leq N^* := \frac{1}{2} \frac{\binom{d}{s}}{\binom{d}{\frac{s}{2}}} \frac{N^s}{(N+1)^{s/2}}.$$

Thus, as long as $|\mathcal{A}| \leq N^*$, there must exist some element $u \in \mathfrak{S}$ such that $\rho_H(u, \mathcal{A}) > \frac{s}{2}$, in which case we can form the augmented set $\mathcal{A} \cup \{u\}$. Iterating this procedure, we can form a set with N^* elements such that $\rho_H(u, v) \geq \frac{s}{2}$ for all $u, v \in \mathcal{A}$.

Finally, we lower bound N^* . We have

$$\begin{aligned} N^* &\stackrel{(i)}{\geq} \frac{1}{2} \left(\frac{d-s}{s/2} \right)^{\frac{s}{2}} \frac{(N)^s}{(N+1)^{s/2}} \\ &= \frac{1}{2} \left(\frac{d-s}{s/2} \right)^{\frac{s}{2}} N^{s/2} \left(\frac{N}{N+1} \right)^{s/2} \\ &\geq \frac{1}{2} \left(\frac{d-s}{s/2} \right)^{\frac{s}{2}} N^{s/2}, \end{aligned}$$

where inequality (i) follows by elementary combinatorics (see Lemma 5 in Raskutti et al., 2011 for details). We conclude that for $s \leq d/4$, we have

$$\log N^* = \Omega\left(s \log \frac{d}{s} + s \log M\left(\frac{\delta}{\sqrt{s}}; \mathbb{B}_{\mathcal{H}}(1), \|\cdot\|_2\right)\right),$$

thereby completing the proof of Lemma 4(a).

E.2 Proof of Part (b)

In order to prove part (b), we instead let $N = M(\frac{1}{2}; \mathbb{B}_{\mathcal{H}}(1), \|\cdot\|_2) - 1$, and then follow the same steps. Since $\log N = \Omega(m)$, we have the modified lower bound

$$\log N^* = \Omega\left(s \log \frac{d}{s} + sm\right),$$

Moreover, instead of the lower bound (51), we have

$$\|g^u - h^v\|_2^2 = \sum_{j=1}^d \|f_j^{u_j} - f_j^{v_j}\|_2^2 \geq \frac{1}{4} \sum_{j=1}^d \mathbb{I}[u_j \neq v_j] \geq \frac{s}{8},$$

using our previous result on the Hamming separation. Furthermore, since $\|f_j\|_2 \leq \|f_j\|_{\mathcal{H}}$ for any univariate function, we have the upper bound

$$\|g^u - h^v\|_2^2 = \sum_{j=1}^d \|f_j^{u_j} - f_j^{v_j}\|_2^2 \leq \sum_{j=1}^d \|f_j^{u_j} - f_j^{v_j}\|_{\mathcal{H}}^2.$$

By the definition (50) of \mathfrak{S} , at most $2s$ of the terms $f_j^{u_j} - f_j^{v_j}$ can be non-zero. Moreover, by construction we have $\|f_j^{u_j} - f_j^{v_j}\|_{\mathcal{H}} \leq 2$, and hence

$$\|g^u - h^v\|_2^2 \leq 8s.$$

Finally, by rescaling the functions by $\sqrt{8}\delta/\sqrt{s}$, we obtain a class of N^* rescaled functions $\{\tilde{g}^u, u \in I\}$ such that

$$\|\tilde{g}^u - \tilde{h}^v\|_2^2 \geq \delta^2, \quad \text{and} \quad \|\tilde{g}^u - \tilde{h}^v\|_2^2 \leq 64\delta^2,$$

as claimed.

Appendix F. Results For Proof Of Theorem 3

The reader should recall from Section 3.5 the definitions of the function classes $\mathcal{F}_{d,s,\mathcal{H}}^*(B)$ and $\mathcal{H}(S,B)$. The function class $\mathcal{H}(S,B)$ can be parameterized by the two-dimensional sequence $(a_{j,k})_{j \in S, k \in \mathbb{N}}$ of co-efficients, and expressed in terms of two-dimensional sequence of basis functions $(\phi_{j,k})_{j \in S, k \in \mathbb{N}}$ and the sequence of eigenvalues $(\mu_k)_{k \in \mathbb{N}}$ for the univariate RKHS \mathcal{H} as follows:

$$\mathcal{H}(S,B) := \left\{ f = \sum_{j \in S} \sum_{k=1}^{\infty} a_{j,k} \phi_{j,k} \mid \sum_{k=1}^{\infty} \frac{a_{j,k}^2}{\mu_k} \leq 1 \forall j \in S \text{ and } \|f\|_{\infty} \leq B \right\}.$$

For any integer $M \geq 1$, we also consider the truncated function class

$$\mathcal{H}(S,B,M) := \left\{ f = \sum_{j \in S} \sum_{k=1}^M a_{j,k} \phi_{j,k} \mid \sum_{k=1}^{\infty} \frac{a_{j,k}^2}{\mu_k} \leq 1 \forall j \in S \text{ and } \|f\|_{\infty} \leq B \right\}.$$

Lemma 10 *We have the inclusion $\mathcal{H}(S, B, M) \subseteq \{f \in \mathcal{H}(S) \mid \sum_{j \in S} \sum_{k=1}^M |a_{j,k}| \leq B\sqrt{M}\}$.*

Proof Without loss of generality, let us assume that $S = \{1, 2, \dots, s\}$, and consider a function $f = \sum_{j=1}^s f_j \in \mathcal{H}(S, B, M)$. Since each f_j acts on a different co-ordinate, we are guaranteed that $\|f\|_\infty = \sum_{j=1}^s \|f_j\|_\infty$. Consider any univariate function $f_j = \sum_{k=1}^M a_{j,k} \phi_{j,k}$. We have

$$\sum_{k=1}^M |a_{j,k}| \leq \sqrt{M} \left(\sum_{k=1}^M a_{j,k}^2 \right)^{1/2} \stackrel{(a)}{\leq} \sqrt{M} [\mathbb{E}[f_j^2(X_j)]]^{1/2} \leq \sqrt{M} \|f_j\|_\infty,$$

where step (a) uses the fact that $\mathbb{E}[f_j^2(X_j)] = \sum_{k=1}^\infty a_{j,k}^2 \geq \sum_{k=1}^M a_{j,k}^2$ for any $M \geq 1$. Adding up the bounds over all co-ordinates, we obtain

$$\|a\|_1 = \sum_{j=1}^s \sum_{k=1}^M |a_{j,k}| \leq \sqrt{M} \sum_{j=1}^s \|f_j\|_\infty = \sqrt{M} \|f\|_\infty \leq \sqrt{MB},$$

where the final step uses the uniform boundedness condition. ■

F.1 Proof of Lemma 5

Recalling the definition of $\widehat{Z}_n(w; t, \mathcal{H}(S, 2B))$ stated from (32), let us view it as a function of the standard Gaussian random vector (w_1, \dots, w_n) . It is straightforward to verify that this variable is Lipschitz (with respect to the Euclidean norm) with parameter at most t/\sqrt{n} . Consequently, by concentration for Lipschitz functions (see Ledoux, 2001), we have

$$\mathbb{P}[\widehat{Z}_n(w; t, \mathcal{H}(S, 2B)) \geq \mathbb{E}[\widehat{Z}_n(w; t, \mathcal{H}(S, 2B))] + 3t\delta_n] \leq \exp\left(-\frac{9n\delta_n^2}{2}\right).$$

Next we prove an upper bound on the expectations

$$\begin{aligned} \widehat{Q}_{w,n}(t; \mathcal{H}(S, 2B)) &:= \mathbb{E}_w \left[\sup_{\substack{g \in \mathcal{H}(S, 2B) \\ \|g\|_n \leq t}} \frac{1}{n} \sum_{i=1}^n w_i g(x_i) \right], \quad \text{and} \\ Q_{w,n}(t; \mathcal{H}(S, 2B)) &:= \mathbb{E}_{x,w} \left[\sup_{\substack{g \in \mathcal{H}(S, 2B) \\ \|g\|_2 \leq t}} \frac{1}{n} \sum_{i=1}^n w_i g(x_i) \right]. \end{aligned}$$

Lemma 11 *Under the conditions of Theorem 3, we have*

$$\max \left\{ \widehat{Q}_{w,n}(t; \mathcal{H}(S, 2B)), Q_{w,n}(t; \mathcal{H}(S, 2B)) \right\} \leq 8BC \sqrt{\frac{s^{1/\alpha} \log s}{n}}.$$

Proof By definition, any function $g \in \mathcal{H}(S, 2B)$ has support at most $2s$, and without loss of generality (re-indexing as necessary), we assume that $S = \{1, 2, \dots, 2s\}$. We can thus view functions in $\mathcal{H}(S, 2B)$ as having domain \mathbb{R}^{2s} , and we can an operator Φ that maps from \mathbb{R}^{2s} to $[\ell^2(\mathbb{N})]^{2s}$, via

$$x \mapsto \Phi_{j,k}(x) = \phi_{j,k}(x_j), \quad \text{for } j = 1, \dots, 2s, \text{ and } k \in \mathbb{N}.$$

Any function in $g \in \mathcal{H}(S, 2B)$ can be expressed in terms of two-dimensional sequence $(a_{j,k})$ and the functions $(\Phi_{j,k})$ as $g(x) = g(x_1, x_2, \dots, x_{2s}) = \sum_{j=1}^{2s} \sum_{k=1}^{\infty} \Phi_{j,k}(x) a_{j,k} = \langle \langle \Phi(x), a \rangle \rangle$, where $\langle \langle \cdot, \cdot \rangle \rangle$ is a convenient shorthand for the inner product between the two arrays.

For any function $g \in \mathcal{H}(S, 2B)$, triangle inequality yields the upper bound

$$\sup_{g \in 2\mathcal{H}(S, 2B)} \frac{1}{n} \left| \sum_{i=1}^n w_i \langle \langle \Phi(x_i), a \rangle \rangle \right| \leq \underbrace{\sup_{g \in 2\mathcal{H}(S, 2B)} \frac{1}{n} \left| \sum_{i=1}^n w_i \langle \langle \Phi_{\cdot, 1:M}(x_i), a_{\cdot, 1:M} \rangle \rangle \right|}_{A_1} + A_2 \quad (52)$$

where $A_2 := \sup_{g \in 2\mathcal{H}(S, 2B)} \frac{1}{n} \left| \sum_{i=1}^n w_i \langle \langle \Phi_{\cdot, M+1:\infty}(x_i), a_{\cdot, M+1:\infty} \rangle \rangle \right|$.

F.1.1 BOUNDING THE QUANTITIES $\mathbb{E}_{x,w}[A_1]$ AND $\mathbb{E}_w[A_1]$

By Hölder's inequality and Lemma 10, we have

$$A_1 \leq \frac{1}{\sqrt{n}} \sup_{g \in 2\mathcal{H}(S, 2B)} \|a_{\cdot, 1:M}\|_{1,1} \max_{j,k} \left| \sum_{i=1}^n \frac{w_i}{\sqrt{n}} \Phi_{j,k}(x_i) \right| \leq \frac{2B\sqrt{M}}{\sqrt{n}} \max_{j,k} \left| \sum_{i=1}^n \frac{w_i}{\sqrt{n}} \Phi_{j,k}(x_i) \right|.$$

By assumption, we have $|\Phi_{j,k}(x_i)| \leq C$ for all indices (i, j, k) , implying that $\sum_{i=1}^n \frac{w_i}{\sqrt{n}} \Phi_{j,k}(x_i)$ is zero-mean with sub-Gaussian parameter bounded by C and we are taking the maximum of $2s \times M$ such terms. Consequently, we conclude that

$$\mathbb{E}_w[A_1] \leq 4BC \sqrt{\frac{M \log(2Ms)}{n}}. \quad (53)$$

Note that the same bound holds for $\mathbb{E}_{x,w}[A_1]$.

F.1.2 BOUNDING THE QUANTITIES $\mathbb{E}_{x,w}[A_2]$ AND $\mathbb{E}_w[A_2]$

In order to control this term, we simply recognize that it corresponds to the usual Gaussian complexity of the sum of $2s$ univariate Hilbert spaces, each of which is an RKHS truncated to the eigenfunctions $\{\mu_k\}_{k \geq M+1}$. In particular, we have

$$\begin{aligned} \frac{1}{n} \left| \sum_{i=1}^n w_i \langle \langle \Phi_{\cdot, M+1:\infty}(x_i), a_{\cdot, M+1:\infty} \rangle \rangle \right| &\leq \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \left| \sum_{k \geq M+1} a_{j,k} \underbrace{\sum_{i=1}^n \Phi_{j,k}(x_i) \frac{w_i}{\sqrt{n}}}_{b_{j,k}} \right| \\ &= \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \left| \sum_{k \geq M+1} \frac{a_{j,k}}{\sqrt{\mu_k}} \sqrt{\mu_k} b_{j,k} \right| \\ &\stackrel{(i)}{\leq} \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \sqrt{\sum_{k \geq M+1} \frac{a_{j,k}^2}{\mu_k}} \sqrt{\sum_{k \geq M+1} \mu_k b_{j,k}^2} \\ &\stackrel{(ii)}{\leq} \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \sqrt{\sum_{k \geq M+1} \mu_k b_{j,k}^2}, \end{aligned}$$

where step (i) follows by applying the Cauchy-Schwarz inequality, and step (ii) exploits the fact that $\sum_{k \geq M+1} \frac{a_{j,k}^2}{\mu_k} \leq 1$ for all j .

This bound no longer depends on the coefficients a (or equivalently, the function g), so that we have shown that

$$\mathbb{E}_w[A_2] \leq \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \mathbb{E}_w \left[\sqrt{\sum_{k \geq M+1} \mu_k b_{j,k}^2} \right] \leq \frac{1}{\sqrt{n}} \sum_{j=1}^{2s} \sqrt{\sum_{k \geq M+1} \mu_k \mathbb{E}_w[b_{j,k}^2]},$$

where the second step uses Jensen's inequality to move the expectation inside the square root. Recalling that $b_{j,k}^2 = \left(\sum_{i=1}^n \Phi_{j,k}(x_i) \frac{w_i}{\sqrt{n}} \right)^2$ and using the independence of the noise variable $\{w_i\}_{i=1}^n$, we have

$$\mathbb{E}_w[b_{j,k}^2] = \frac{1}{n} \sum_{i=1}^n \Phi_{j,k}^2(x_i) \mathbb{E}_w[w_i^2] \leq C^2.$$

Putting together the pieces, we conclude that

$$\mathbb{E}_w[A_2] \leq \frac{C}{\sqrt{n}} \sum_{j=1}^{2s} \sqrt{\sum_{k \geq M+1} \mu_k} = \frac{2Cs}{\sqrt{n}} \sqrt{\sum_{k \geq M+1} \mu_k}. \tag{54}$$

Once again, a similar bound holds for $\mathbb{E}_{x,w}[A_2]$.

Substituting the bounds (53) and (54) into the inequality (52), we conclude that

$$\begin{aligned} Q_{w,n}(2\mathcal{H}(S, 2B)) &\leq 4BC \sqrt{\frac{M \log(2Ms)}{n}} + 2Cs \sqrt{\frac{\sum_{k \geq M+1} \mu_k}{n}} \\ &\leq 4BC \sqrt{\frac{M \log(2Ms)}{n}} + 2Cs \sqrt{\frac{M^{1-2\alpha}}{n}}, \end{aligned}$$

where the second inequality follows from the relation $\mu_k \simeq k^{-2\alpha}$. Finally, setting $M = s^{\frac{1}{\alpha}}$ yields the claim. Note that the same argument works for the Rademacher complexity, since we only exploited the sub-Gaussianity of the variables w_i . This completes the proof of Lemma 11. \blacksquare

Returning to the proof of Lemma 5, combining Lemma 11 with the bound (40) in Lemma 8:

$$\mathbb{P}[\widehat{Z}_n(w; t, \mathcal{H}(S, 2B)) \geq 8BC \sqrt{\frac{s^{1/\alpha} \log s}{n}} + 3t\delta_n] \leq \exp\left(-\frac{9n\delta_n^2}{2}\right).$$

Since $\|g\|_n \leq 2B$ for any function $g \in \mathcal{H}(S, 2B)$, the proof Lemma 5 is completed using a peeling argument over the radius, analogous to the proof of Lemma 1 (see Appendix B).

References

- K. S. Alexander. Rates of Growth and Sample Moduli for Weighted Empirical Processes Indexed by Sets. *Probability Theory and Related Fields*, 75:379–423, 1987.
- N. Aronszajn. Theory of Reproducing Kernels. *Transactions of the American Mathematical Society*, 68:337–404, 1950.
- F. Bach. Consistency of the Group Lasso and Multiple Kernel Learning. *Journal of Machine Learning Research*, 9:1179–1225, 2008.

- P. Bartlett, O. Bousquet, and S. Mendelson. Local Rademacher Complexities. *Annals of Statistics*, 33:1497–1537, 2005.
- P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous Analysis of Lasso and Dantzig Selector. *Annals of Statistics*, 37(4):1705–1732, 2009.
- M. S. Birman and M. Z. Solomjak. Piecewise-polynomial Approximations of Functions of the Classes W_p^α . *Math. USSR-Sbornik*, 2(3):295–317, 1967.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- L. Breiman. Better Subset Regression Using the Nonnegative Garrote. *Technometrics*, 37:373–384, 1995.
- V. V. Buldygin and Y. V. Kozachenko. *Metric Characterization of Random Variables and Random Processes*. American Mathematical Society, Providence, RI, 2000.
- B. Carl and I. Stephani. *Entropy, Compactness and the Approximation of Operators*. Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge, UK, 1990.
- B. Carl and H. Triebel. Inequalities Between Eigenvalues, Entropy Numbers and Related Quantities of Compact Operators in Banach Spaces. *Annals of Mathematics*, 251:129–133, 1980.
- F. Chung and L. Lu. Concentration Inequalities and Martingale Inequalities. *Internet Mathematics*, 3:79–127, 2006.
- T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, 1991.
- U. Einmahl and D. M. Mason. Some Universal Results on the Behavior of the Increments of Partial Sums. *Annals of Probability*, 24:1388–1407, 1996.
- C. Gu. *Smoothing Spline ANOVA Models*. Springer Series in Statistics. Springer, New York, NY, 2002.
- R. Z. Has'minskii. A Lower Bound on the Risks of Nonparametric Estimates of Densities in the Uniform Metric. *Theory Prob. Appl.*, 23:794–798, 1978.
- T. Hastie and R. Tibshirani. Generalized Additive Models. *Statistical Science*, 1(3):297–310, 1986.
- G. Kimeldorf and G. Wahba. Some Results on Tchebycheffian Spline Functions. *Jour. Math. Anal. Appl.*, 33:82–95, 1971.
- V. Koltchinskii and M. Yuan. Sparse Recovery in Large Ensembles of Kernel Machines. In *Proceedings of COLT*, 2008.
- V. Koltchinskii and M. Yuan. Sparsity in Multiple Kernel Learning. *Annals of Statistics*, 38:3660–3695, 2010.
- M. Ledoux. *The Concentration of Measure Phenomenon*. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2001.

- Y. Lin and H. H. Zhang. Component Selection and Smoothing in Multivariate Nonparametric Regression. *Annals of Statistics*, 34:2272–2297, 2006.
- P. Massart. About the Constants in Talagrand’s Concentration Inequalities for Empirical Processes. *Annals of Probability*, 28(2):863–884, 2000.
- L. Meier, S. van de Geer, and P. Bühlmann. High-dimensional Additive Modeling. *Annals of Statistics*, 37:3779–3821, 2009.
- S. Mendelson. Geometric Parameters of Kernel Machines. In *Proceedings of COLT*, pages 29–43, 2002.
- J. Mercer. Functions of Positive and Negative Type and Their Connection With the Theory of Integral Equations. *Philosophical Transactions of the Royal Society A*, 209:415–446, 1909.
- S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A Unified Framework for High-dimensional Analysis of M -estimators with Decomposable Regularizers. In *NIPS Conference*, 2009.
- G. Pisier. *The Volume of Convex Bodies and Banach Space Geometry*, volume 94 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, UK, 1989.
- G. Raskutti, M. J. Wainwright, and B. Yu. Minimax Rates of Estimation for High-dimensional Linear Regression Over ℓ_q -balls. *IEEE Trans. Information Theory*, 57(10):6976–6994, October 2011.
- P. Ravikumar, H. Liu, J. Lafferty, and L. Wasserman. SpAM: Sparse Additive Models. *Journal of the Royal Statistical Society, Series B*, 71(5):1009–1030, 2009.
- S. Saitoh. *Theory of Reproducing Kernels and its Applications*. Longman Scientific & Technical, Harlow, UK, 1988.
- B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- C. J. Stone. Additive Regression and Other Nonparametric Models. *Annals of Statistics*, 13(2): 689–705, 1985.
- T. Suzuki and M. Sugiyama. Fast Learning Rate of Multiple Kernel Learning: Trade-off Between Sparsity and Smoothness. In *AISTATS Conference*, 2012.
- S. van de Geer. *Empirical Processes in M -Estimation*. Cambridge University Press, 2000.
- A. W. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, New York, NY, 1996.
- G. Wahba. *Spline Models for Observational Data*. CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, PA, 1990.
- Y. Yang and A. Barron. Information-theoretic Determination of Minimax Rates of Convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.

B. Yu. Assouad, Fano and Le Cam. *Research Papers in Probability and Statistics: Festschrift in Honor of Lucien Le Cam*, pages 423–435, 1996.

M. Yuan. Nonnegative Garrote Component Selection in Functional ANOVA Models. In *Conference on Artificial Intelligence and Statistics*, pages 660–666, 2007.