

RISK BOUNDS FOR HIGH-DIMENSIONAL RIDGE FUNCTION COMBINATIONS INCLUDING NEURAL NETWORKS

BY JASON M. KLUSOWSKI AND ANDREW R. BARRON

Yale University

Let f^* be a function on \mathbb{R}^d satisfying a spectral norm condition. For various noise settings, we show that $\mathbb{E}\|\hat{f} - f^*\|^2 \leq v_{f^*} \left(\frac{\log d}{n}\right)^{1/4}$, where n is the sample size and \hat{f} is either a penalized least squares estimator or a greedily obtained version of such using linear combinations of ramp, sinusoidal, sigmoidal or other bounded Lipschitz ridge functions. Our risk bound is effective even when the dimension d is much larger than the available sample size. For settings where the dimension is larger than the square root of the sample size this quantity is seen to improve the more familiar risk bound of $v_{f^*} \left(\frac{d \log(n/d)}{n}\right)^{1/2}$, also investigated here.

1. Introduction. Functions f^* in \mathbb{R}^d are approximated using linear combinations of ridge functions with one layer of nonlinearities. These approximations are employed via functions of the form

$$(1.1) \quad f_m(x) = f_m(x, \zeta) = \sum_{k=1}^m c_k \phi(a_k \cdot x + b_k),$$

which is parameterized by the vector ζ , consisting of a_k in \mathbb{R}^d , and b_k, c_k in \mathbb{R} for $k = 1, \dots, m$, where $m \geq 1$ is the number of nonlinear terms. The functions ϕ are allowed to be quite general. For example, they can be bounded and Lipschitz, polynomials with certain controls on their degrees, or bounded with jump discontinuities. It is useful to view the representation (1.1) as

$$(1.2) \quad \sum_{h \in \mathcal{H}} \beta_h h(x),$$

Primary 62J02, 62G08; secondary 68T05

Keywords and phrases: Nonparametric regression; nonlinear regression; neural networks; penalization; machine learning; high dimensional data analysis; learning theory; generalization error; supervised learning; greedy algorithms

where \mathcal{H} is a library of candidate basis or activation functions of the form $h(x) = \phi(\theta_h \cdot x)$ and the β_h are non-negative, where all but finitely many are zero.

We can reduce (1.1) to (1.2) as follows. Suppose the library is symmetric $\mathcal{H} = -\mathcal{H}$ and contains the zero function. Without loss of generality, we may assume that the c_k are non-negative by replacing the associated ϕ with $\phi \operatorname{sgn} c_k$, that by assumption also belongs to \mathcal{H} . One can assume the internal parameters $a_k \cdot x + b_k$ take the form $\theta_k \cdot x$ by appending a one to x and b_k to a_k . Note that now x and θ_k are $(d+1)$ -dimensional.

Suppose P is an arbitrary probability measure on $[-1, 1]^d$. Let $\|\cdot\|$ be the $L^2(P)$ norm induced by the inner product $\langle \cdot, \cdot \rangle$. Define $v_{f^*,s} = \int_{\mathbb{R}^d} \|\omega\|_1^s \tilde{f}(\omega) d\omega$, for $s \geq 0$. If f^* has a bounded domain in $[-1, 1]^d$ and a Fourier representation $f^*(x) = \int_{\mathbb{R}^d} e^{i\omega \cdot x} \tilde{f}(\omega) d\omega$ with $v_{f^*,1} < +\infty$, it is possible to use approximating functions of the form (1.1) with a single activation function ϕ . Such activation functions ϕ can be general bounded monotone functions.

The following result from [2] provides a useful starting point for approximating general functions f^* by linear combinations of such objects. Suppose $v_{f^*,1}$ is finite. Then by [2], there exists an artificial neural network of the form (1.1) with $\phi(x) = \operatorname{sgn}(x)$ with $\|a_k\|_1 = 1$ and $|b_k| \leq 1$ such that

$$\|f^* - f_m\|^2 \leq \frac{4v_{f^*,1}^2}{m}.$$

If ϕ has right at left limits -1 and $+1$, respectively, the fact that $\phi(\tau x) \rightarrow \operatorname{sgn}(x)$ as $\tau \rightarrow +\infty$ allows one to use somewhat arbitrary activation functions as basis elements. It is natural to impose a restriction on the size of the internal parameters and to also enjoy a certain degree of smoothness not offered by step functions. Thus we consider the result in [10], which allows one to approximate f^* by linear combinations of ramp ridge functions (also known as first order ridge splines or hinging hyper-planes) $(x \cdot \alpha - t)_+ = \max\{0, x \cdot \alpha - t\}$, with $\|\alpha\|_1 = 1$, $|t| \leq 1$. These functions are continuous and Lipschitz. In the supplementary material [Supplement A](#) we refine a result from [10]. For an arbitrary target function f^* with $v_{f^*,2}$ finite, there exists an approximation of the form (1.1) activated by ridge ramp functions with $\|a_k\| = 1$ and $|b_k| \leq 1$ such that

$$\|f^* - f_m\|^2 \leq \frac{16v_{f^*,2}^2}{m}.$$

The supplement also discusses refinements of these approximation bounds and how one can reach similar conclusions with second order splines having

bounded internal parameters. The Lipschitz property of the ramp functions yields smaller covering numbers and thus improved rates over that which can be obtained using step functions. We define the set

$$\mathcal{H}_{\text{ramp}} = \{x \mapsto \pm(\alpha \cdot x - t)_+ : \|\alpha\|_1 = 1, |t| \leq 1\}.$$

We then set $\mathcal{F}_{\text{ramp}}$ to be the linear span of $\mathcal{H}_{\text{ramp}}$. In general, for a symmetric collection of dictionary elements $\mathcal{H} = -\mathcal{H}$ containing the zero function, we let $\mathcal{F} = \mathcal{F}_{\mathcal{H}}$ be the linear span of \mathcal{H} . The variation $v_f = \|f\|_{\mathcal{H}}$ of f with respect to \mathcal{H} (or the atomic norm of f with respect to \mathcal{H}) is defined by

$$\liminf_{\delta \downarrow 0} \inf_{f_\delta \in \mathcal{F}} \left\{ \|\beta\|_1 : f_\delta = \sum_{h \in \mathcal{H}} \beta_h h \text{ and } \|f_\delta - f\| \leq \delta, \beta_h \in \mathbb{R}^+ \right\},$$

where $\|\beta\|_1 = \sum_{h \in \mathcal{H}} \beta_h$. For functions in $\mathcal{F}_{\mathcal{H}}$, this variation picks out the smallest $\|\beta\|_1$ among representations $f = \sum_{h \in \mathcal{H}} \beta_h h$. For functions in the $L^2(P)$ closure of the linear span of \mathcal{H} , the variation is the smallest limit of such ℓ_1 norms among functions approaching the target. The subspace of functions with $\|f\|_{\mathcal{H}}$ finite is denoted $L_{1,\mathcal{H}}$.

Note that the condition $\int_{\mathbb{R}^d} \|\omega\|_1^2 |\tilde{f}(\omega)| d\omega < +\infty$ ensures that f^* belongs to $L_{1,\mathcal{H}_{\text{ramp}}}$ and $\|f^*\|_{\mathcal{H}_{\text{ramp}}} \leq v_{f^*,2}$. Functions with moderate variation are particularly closely approximated. Nevertheless, even when $\|f^*\|_{\mathcal{H}}$ is infinite, we express the trade-offs in approximation accuracy for consistently estimating functions in the closure of the linear span of \mathcal{H} .

In what follows, we assume h has L_∞ norm at most one, h is Lipschitz with Lipschitz constant at most one, and the internal parameters have ℓ_1 norm at most Λ . This control on the size of the internal parameters will be featured prominently throughout. In the case of ramp activation functions, we are content with the assumption $\Lambda = 2$. Note that if one restricts the size of the domain and internal parameters (say, to handle polynomials), the functions h are still bounded and Lipschitz but with possibly considerably worse constants.

Suppose data $\{(X_i, Y_i)\}_{i=1}^n$ are independently drawn from the distribution of (X, Y) . To produce predictions of the real-valued response Y from its input X , the target regression function $f^*(x) = \mathbb{E}[Y|X = x]$ is to be estimated. The function f^* is assumed to be bounded in magnitude by a positive constant B . We assume the noise $\epsilon = Y - f^*(X)$ has moments (conditioned on X) that satisfy a Bernstein condition with parameter $\eta > 0$. That is, we assume

$$\mathbb{E}[|\epsilon|^k | X] \leq \frac{1}{2} k! \eta^{k-2} \mathbb{V}(\epsilon | X), \quad k = 3, 4, \dots,$$

where $\mathbb{V}(\epsilon|X) \leq \sigma^2$. This assumption is equivalent to requiring that $\mathbb{E}(e^{|\epsilon|/\nu}|X)$ is uniformly bounded in X for some $\nu > 0$. A stricter assumption is that $\mathbb{E}(e^{|\epsilon|^2/\nu}|X)$ is uniformly bounded in X , which corresponds to an error distribution with sub-Gaussian tails. These two noise settings will give rise to different risk bounds, as we will see.

The input design X is assumed to be a d -dimensional vector with sub-Gaussian coordinates. This condition implies that $\|X\|_\infty^2$ is on average bounded by a constant multiple of $\log d$. We will see that in our framework, estimators of f^* are functions of training and test data. If one seeks to describe the error of an estimator for f^* evaluated at a new set of random points distributed according to P , the estimator is allowed to depend on the test data and the sub-Gaussian coordinates can be shown to affect the rates below only by a logarithmic factor in d . This paper will focus only on the ability of an estimator to generalize to a new data set. For ease of analysis, we assume that X is contained in the hyper-cube $[-1, 1]^d$, i.e. the supremum norm of X is at most one. No assumption is made about whether the coordinates of X are independent.

Because f^* is bounded in magnitude by B , it is useful to truncate an estimator \hat{f} at a level B_n at least B . Depending on the nature of the noise ϵ , we will see that B_n will need to be at least B plus a term of order $\sqrt{\log n}$ or $\log n$. We define the truncation operator T that acts on function f in \mathcal{F} by $Tf = \min\{|f|, B_n\} \text{sgn} f$. Associated with the truncation operator is a tail quantity

$$T_n = 2 \sum_{i=1}^n (|Y_i|^2 - B_n^2) \mathbb{I}\{|Y_i| > B_n\}$$

that appears in the following analysis. Lemma 9 describes the finite sample behavior of $\mathbb{E}T_n$.

The empirical mean squared error of a function f as a candidate fit to the observed data is $(1/n) \sum_{i=1}^n (Y_i - f(X_i))^2$. Given the collection of functions \mathcal{F} , a penalty $\text{pen}_n(f)$, $f \in \mathcal{F}$, and data, a penalized least squares estimator \hat{f} arises by optimizing or approximately optimizing

$$(1/n) \sum_{i=1}^n (Y_i - f(X_i))^2 + \text{pen}_n(f)/n.$$

Our method of risk analysis proceeds as follows. Given a collection \mathcal{F} of candidate functions, we show that there is a countable approximating set $\tilde{\mathcal{F}}$ of representations \tilde{f} , variable-distortion, variable-complexity cover of \mathcal{F} , and

a complexity function $L_n(\tilde{f})$, with the property that for each f in \mathcal{F} , there is an \tilde{f} in $\tilde{\mathcal{F}}$ such that $\text{pen}_n(f)$ is not less than a constant multiple of $\gamma_n L_n(\tilde{f}) + \Delta_n(f, \tilde{f})$, where γ_n is a constant (depending on B , B_n , σ^2 , and η) and $\Delta_n(f, \tilde{f})$ is given as a suitable empirical measure of distortion (based on sums of squared errors). The variable-distortion, variable-complexity terminology has its origins in [6], [7], and [12]. The task is to determine penalties such that an estimator \hat{f} approximately achieving the minimum of $\|Y - f\|_n^2 + \text{pen}_n(f)/n$ satisfies

$$(1.3) \quad \mathbb{E}\|T\hat{f} - f^*\|^2 \leq c \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n + \mathbb{E}A_f\},$$

for some universal $c > 1$ and positive quantity A_f that decays as n grows. The quantity

$$\inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n + \mathbb{E}A_f\}.$$

is an index of resolvability of f^* by functions \mathcal{F} with sample size n . We shall take particular advantage of such risk bounds in the case that $\text{pen}_n(f)$ does not depend on \underline{X} . Our restriction of X to $[-1, 1]^d$ is one way to allow the construction of such penalties.

The following table expresses the heart of our results, expressing valid penalties providing such risk bounds for moderate and high-dimensional situations.

Noise ϵ	$\lambda_n \gtrsim$	$\text{pen}_n(f)/n \gtrsim$
sub-Gaussian / sub-exponential	$\left(\frac{\gamma_n^2 \log(d+1)}{n}\right)^{1/4}$	$v_f \lambda_n$
zero	$\left(\frac{\gamma_n \log(d+1)}{n}\right)^{1/3}$	$(v_f)^{4/3} \lambda_n$
sub-Gaussian / sub-exponential	$\left(\frac{d\gamma_n \log(n/d+1)}{n}\right)^{1/2+1/(2(d+3))}$	$v_f \lambda_n$

Table 1: Penalties for Theorem 2

The results we wish to highlight are contained in the first two rows of Table 1. The penalties as stated are valid up to modest universal constants and negligible terms that do not depend on the candidate fit. The quantity γ_n is of order $\log^2 n$ in the sub-exponential noise case, order $\log n$ in the sub-Gaussian noise case and of constant order in the zero noise case. This γ_n (as defined in Lemma 9) depends on the variance bound σ^2 , Bernstein parameter η , the upper bound B of $\|f^*\|_{\mathcal{H}}$, and the noise tail level B_n of the indicated order.

When f^* belongs to $L_{1,\mathcal{H}}$, a resulting valid risk bound is a constant multiple of $\|f^*\|_{\mathcal{H}}\lambda_n$ or $\|f^*\|_{\mathcal{H}}^{4/3}\lambda_n$, according to the indicated cases. In this way the λ_n expression provides a rate of convergence.

The classical risk bounds for mean squared error, involving d/n to some power, are only useful when the sample size is much larger than the dimension. Here, in contrast, in the first two lines of Table 1, we see the dependence on dimension is logarithmic, permitting much smaller sample sizes. The price we pay for the smaller dependence on dimension is a deteriorated rate with exponent $1/4$ in general and $1/3$ under a no noise assumption. The rates in last row improve upon the familiar exponent of $1/2$ to $1/2 + 1/(2(d+3))$. Note that when d is large, this enhancement in the exponent is negligible. The rate in the first row is better than the third approximately for $d > \sqrt{n}$, the second is better than the third row approximately for $d > n^{1/3}$, and both of these first two rows have risk tending to zero as long as $d < e^{o(n)}$.

For functions in $L_{1,\mathcal{H}_{\text{ramp}}}$, an upper bound of $((d/n)\log(n/d))^{1/2}$ for the squared error loss is obtained in [3]. Using the truncated penalized ℓ_1 least squares estimator (1.3), we obtain an improved rate of order $((d\gamma_n/n)\log(n/d))^{1/2+1/(2(d+3))}$, where γ_n is logarithmic in n , using techniques that originate in [15] and [14], with some corrections here. A slightly better rate with the $d+3$ replaced by $d+1$ in the denominator of the exponent can be achieved through more technical means, that we choose not to put in the present paper to keep the length under control.

In an upcoming paper, the authors intend to show that this rate is almost optimal in the Gaussian noise, uniform design setting, since we provide mini-max rates between $(1/n)^{1/2+1/(d+2)}$ and $(d/n)^{1/2+1/2(d+1)}$ for functions in $L_{1,\mathcal{H}_{\text{ramp}}}$. Compare this with [20], where the mini-max L^2 risk for functions in $L_{1,\mathcal{H}_{\text{step}}}$ (i.e. function approximated by linear combinations of step ridge functions) is determined to be between $(1/n)^{1/2+1/(2(d+1))}(\log n)^{-(1+1/d)(1+2/d)(1+2/d)(2+1/d)}$ and $(\log n/n)^{1/2+1/(2(2d+1))}$.

These quantities have the attractive feature that the rate does not deteriorate as the dimension grows. However, they are only useful provided d/n is small. In high dimensional settings, the available sample size might not be large enough to ensure this condition. These results are all based on obtaining covering numbers for the library $\{x \mapsto \phi(\theta \cdot x) : \|\theta\|_1 \leq \Lambda\}$. If ϕ satisfies a Lipschitz condition, these numbers are equivalent to ℓ_1 covering numbers of the internal parameters or of the Euclidean inner product of the data and the internal parameters. The factor of d multiplying the reciprocal of the sample size is produced from the order $d\log(\Lambda/\epsilon)$ log cardinality of

the standard covering of the library $\{\theta : \|\theta\|_1 \leq \Lambda\}$. What enables us to circumvent this polynomial dependence on d is to use an alternative cover of $\{x \mapsto x \cdot \theta : \|\theta\|_1 \leq \Lambda\}$ that has log cardinality of order $(\Lambda/\epsilon)^2 \log(d+1)$. Misclassification errors for neural networks with bounded internal parameters have been analyzed in [9] and [18].

In this paper we bound the mean squared error of function estimation, extending the results of [3], [4], [20], [15], [8], [7], [18], [11], [1], [21], [17].

2. Greedy Algorithm. The main difficulty with constructing an estimator that satisfies (1.3) is that it involves a dm -dimensional optimization. Here, we outline a greedy approach that reduces the problem to performing m d -dimensional optimizations. This construction is based on the ℓ_1 -penalized greedy pursuit (LPGP) in [15], with the modification that the penalty can be a convex function of the candidate function complexity. Greedy strategies for approximating functions in the closure of the linear span of a subset of a Hilbert space has its origins in [16] and many of its statistical implications were studied in [8] and [15].

Let f^* be a function, not necessarily in \mathcal{F} . Initialize $f_0 = 0$. For $m = 1, 2, \dots$, iteratively, given the terms of f_{m-1} as h_1, \dots, h_{m-1} and the coefficients of it as $\beta_{1,m-1}, \dots, \beta_{m-1,m-1}$, we proceed as follows. Let $f_m(x) = \sum_{j=1}^m \beta_{j,m} h_j(x) = \sum_{j=1}^m \beta_{j,m} \phi(\theta_{h_j} \cdot x)$, with the term h_m in \mathcal{H} chosen to come within a constant factor $c \geq 1$ of the maximum inner product with the residual $f^* - f_{m-1}$; that is

$$\langle h_m, f^* - f_{m-1} \rangle \geq \frac{1}{c} \sup_{h \in \mathcal{H}} \langle h, f^* - f_{m-1} \rangle.$$

Define $f_m(x) = (1 - \alpha_m)f_{m-1}(x) + \beta_{m,m}h_m(x)$. Associated with this representation of f_m is the ℓ_1 norm of its coefficients $v_m = \sum_{j=1}^m |\beta_{j,m}| = (1 - \alpha_m)v_{m-1} + \beta_{m,m}$. The coefficients α_m and $\beta_{m,m}$ are chosen so that

$$\begin{aligned} & \|f^* - (1 - \alpha_m)f_{m-1} - \beta_{m,m}h_m\|^2 + w(v_m) \\ \leq & \inf_{\alpha \in [0,1], \beta \in \mathbb{R}^+} [\|f^* - (1 - \alpha)f_{m-1} - \beta h_m\|^2 + w((1 - \alpha)v_{m-1} + \beta)], \end{aligned}$$

where $w : \mathbb{R} \rightarrow \mathbb{R}$ is a real-valued non-negative convex function. In the empirical setting, with $R_i = Y_i - f_{m-1}(X_i)$, the high-dimensional optimization task is to find θ_m such that

$$\frac{1}{n} \sum_{i=1}^n R_i \phi(\theta_m \cdot X_i) \geq \frac{1}{c} \sup_{\|\theta\|_1 \leq \Lambda} \frac{1}{n} \sum_{i=1}^n R_i \phi(\theta \cdot X_i)$$

One appealing aspect of the slight sub-optimality of θ_m is that it permits the use of adaptive annealing techniques developed by the authors in a forthcoming paper [5]. The algorithm samples from a distribution proportional to $e^{t(\frac{1}{n} \sum_{i=1}^n R_i \phi(\theta \cdot X_i))} p_0(\theta)$ which has, for t sufficiently large, a mean that is at least $\frac{1}{c} \sup_{\|\theta\|_1 \leq \Lambda} \frac{1}{n} \sum_{i=1}^n R_i \phi(\theta \cdot X_i)$.

THEOREM 1. *If f_m is chosen according to the greedy scheme described previously, then*

$$\|f^* - f_m\|^2 + w(v_m) \leq \inf_{f \in \mathcal{F}} \left\{ \|f^* - f\|^2 + w(cv_f) + \frac{4b_f}{m} \right\},$$

where $b_f = c^2 v_f^2 + 2v_f \|f^*\| (c+1) - \|f\|^2$. Furthermore,

$$\inf_{f \in \mathcal{F}} \inf_{\delta > 0} \left\{ (1 + \delta) \|f^* - f\|^2 + w(cv_f) + \frac{4(1 + \delta)\delta^{-1}(c+1)^2 v_f^2}{m} \right\},$$

and hence with $\delta = \frac{2(c+1)v_f}{\|f^* - f\|\sqrt{m}}$,

$$\|f^* - f_m\|^2 + w(v_m) \leq \inf_{f \in \mathcal{F}} \left\{ \left(\|f^* - f\| + \frac{2(c+1)v_f}{\sqrt{m}} \right)^2 + w(cv_f) \right\}.$$

PROOF. Fix any f in the linear span \mathcal{F} , with the form $\sum_{h \in \mathcal{H}} \beta_h h$, with non-negative β_h .

$$e_m = \|f^* - f_m\|^2 - \|f^* - f\|^2 + w(v_m).$$

Then from the definition of α_m and $\beta_{m,m}$,

$$\begin{aligned} e_m &= \|f^* - (1 - \alpha_m)f_{m-1} - \beta_{m,m}h_m\|^2 - \|f^* - f\|^2 + \\ &\quad w((1 - \alpha_m)v_{m-1} + \beta_{m,m}) \\ &\leq \|f^* - (1 - \alpha_m)f_{m-1} - \alpha_m cv_f h_m\|^2 - \|f^* - f\|^2 + \\ &\quad w((1 - \alpha_m)v_{m-1} + \alpha_m cv_f) \\ &\leq \|f^* - (1 - \alpha_m)f_{m-1} - \alpha_m cv_f h_m\|^2 - \|f^* - f\|^2 + \\ &\quad (1 - \alpha_m)w(v_{m-1}) + \alpha_m w(cv_f), \end{aligned}$$

where the last line follows from the convexity of w . Now $\|f^* - (1 - \alpha_m)f_{m-1} - \alpha_m cv_f h_m\|^2$ is equal to $\|(1 - \alpha_m)(f^* - f_{m-1}) + \alpha_m(f^* - ch_m v_f)\|^2$. Expanding

this quantity leads to

$$\begin{aligned} \|f^* - (1 - \alpha_m)f_{m-1} - \alpha_m cv_f h_m\|^2 &= (1 - \alpha_m)^2 \|f^* - f_{m-1}\|^2 \\ &\quad - 2\alpha_m(1 - \alpha_m)\langle f^* - f_{m-1}, ch_m v_f - f^* \rangle \\ &\quad + \alpha_m^2 \|f^* - ch_m v_f\|^2. \end{aligned}$$

Next we add $(1 - \alpha_m)w(v_{m-1}) + \alpha_m w(cv_f) - \|f^* - f\|^2$ to this expression to obtain

$$\begin{aligned} e_m &\leq (1 - \alpha_m)e_{m-1} + \alpha_m^2 [\|f^* - ch_m v_f\|^2 - \|f^* - f\|^2] + \alpha_m w(cv_f) \\ &\quad - 2\alpha_m(1 - \alpha_m)\langle f^* - f_{m-1}, ch_m v_f - f \rangle \\ &\quad + \alpha_m(1 - \alpha_m)[2\langle f^* - f_{m-1}, f^* - f \rangle - \|f^* - f_{m-1}\|^2 - \|f^* - f\|^2], \end{aligned}$$

which is further upper bounded by

$$\begin{aligned} e_m &\leq (1 - \alpha_m)e_{m-1} + \alpha_m^2 [\|f^* - ch_m v_f\|^2 - \|f^* - f\|^2] + \alpha_m w(cv_f) \\ &\quad - 2\alpha_m(1 - \alpha_m)\langle f^* - f_{m-1}, ch_m v_f - f \rangle \\ &\quad - \alpha_m(1 - \alpha_m)(\|f^* - f_{m-1}\| - \|f^* - f\|)^2, \end{aligned}$$

Consider a random variable that equals h with probability β_h/v_f having mean f . Since a maximum is at least an average, the choice of h_m implies that $\langle f^* - f_{m-1}, ch_m v_f \rangle$ is at least $\langle f^* - f_{m-1}, f \rangle$. This shows that e_m is no less than $(1 - \alpha_m)e_{m-1} + \alpha_m^2 [\|f^* - ch_m v_f\|^2 - \|f^* - f\|^2] + \alpha_m w(cv_f)$. Expanding the squares in $\|f^* - ch_m v_f\|^2 - \|f^* - f\|^2$ and using the Cauchy-Schwarz inequality yields the bound $\|ch_m v_f\|^2 + 2\|f^*\|(\|f - ch_m v_f\|) - \|f\|^2$. Since $\|h_m\| \leq 1$ and $\|f\| \leq v_f$, we find that $\|f^* - ch_m v_f\|^2 - \|f^* - f\|^2$ is at most $b_f = c^2 v_f^2 + 2v_f \|f^*\|(c + 1) - \|f\|^2$. Hence we have shown that

$$e_1 \leq b_f + w(cv_f)$$

and

$$e_m \leq (1 - \alpha_m)e_{m-1} + \alpha_m^2 b_f + \alpha_m w(cv_f).$$

Choose $\alpha_m = 2/(m + 1)$, $m \geq 2$ and use an inductive argument to establish the claim. The second statement in the theorem follows from similar arguments upon consideration of

$$e_m = \|f^* - f_m\|^2 - (1 + \delta)\|f^* - f\|^2 + w(v_m),$$

together with the inequality $a^2 - (1 + \delta)b^2 \leq (1 + \delta)\delta^{-1}(a - b)^2$. \square

3. Risk bounds. Here we state our main theorem.

THEOREM 2. *Let f^* be a real-valued function on $[-1, 1]^d$ with finite variation v_{f^*} with respect to the library $\mathcal{H} = \{h(x) = \phi(\theta_h \cdot x)\}$. We further assume that ϕ is Lipschitz function with $\|\phi\|_\infty \leq 1$ and $\|\theta_h\| \leq \Lambda$. In the case that f^* belongs to $L_{1, \mathcal{H}_{ramp}}$, ϕ is a ramp function with $\Lambda = 2$. If \hat{f} is chosen to satisfy*

$$\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{f}(X_i))^2 + \text{pen}_n(\hat{f})/n \leq \inf_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n (Y_i - f(X_i))^2 + \text{pen}_n(f)/n \right\},$$

then for the truncated estimator $T\hat{f}$ and for $\text{pen}_n(f)$ depending on v_f as specified below, the risk has the resolvability bound

$$\mathbb{E}\|T\hat{f} - f^*\|^2 \leq 2(\tau + 1) \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n\},$$

The penalty divided by sample size $\text{pen}_n(f)/n$ is at least

$$16v_f \left(\frac{\gamma_n B_n^2 \Lambda^2 \log(d+1)}{n} \right)^{1/4} + 8 \left(\frac{\gamma_n B_n^2 \Lambda^2 \log(d+1)}{n} \right)^{1/2} + \frac{T_n}{n}$$

and

$$60v_f \Lambda \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+1/2(d+3)} + \frac{1}{\Lambda^2} \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+1/2(d+3)} \\ + \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+3/2(d+3)} + \frac{d\gamma_n \log(n/d+1)}{n} + \frac{T_n}{n},$$

when d is small compared to n . In the no noise setting, $\text{pen}_n(f)/n$ is at least

$$16v_f^{4/3} \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{n} \right)^{1/3} + 4(v_f^{4/3} + 1) \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{n} \right)^{2/3},$$

Here $\gamma_n = (2\tau)^{-1}(1 + \delta_1/2)(1 + 2/\delta_1)(B + B_n)^2 + 2(1 + 1/\delta_2)\sigma^2 + 2(B + B_n)\eta$ and $\tau = (1 + \delta_1)(1 + \delta_2)$ for some $\delta_1 > 0$ and $\delta_2 > 0$. Accordingly, if f^* belongs to $L_{1, \mathcal{H}}$, $\mathbb{E}\|T\hat{f} - f^*\|^2$ is not more than a constant multiple of the above penalties with v_f replaced by $\|f^*\|_{\mathcal{H}}$.

If \hat{f}_m is the LPGP estimator from the previous section, then by Theorem 1,

$$\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{f}_m(X_i))^2 + w(v_{\hat{f}_m}) \leq \inf_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n (Y_i - f(X_i))^2 + w(cv_f) + \frac{4b_f}{m} \right\},$$

where b_f is the empirical version of the same quantity in Theorem 1 and hence the risk has the resolvability bound

$$\mathbb{E}\|T\hat{f} - f^*\|^2 \leq 2(\tau + 1) \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}pen_n(cf)/n + 4\mathbb{E}b_f/m\},$$

for a penalty, convex in v_f , $pen_n(f) = nw(v_f)$ as before. If m is chosen to be of order between \sqrt{n} and n so as to make the computational effects negligible, the previously described $L^2(P)$ rates for estimating f^* in $L_{1,\mathcal{H}}$ via the truncated estimator $T\hat{f}_m$ are attainable under the appropriate penalties.

As we have said, where we have $d + 3$ in the denominator in the exponent, it is possible to improve it to a $d + 1$. Such improvements are due to improved covers for the complexity evaluation. These refinements also hold for the LPGP estimator, although the number of iterations m needs to be of slightly higher order than before.

One can also extend these results to include penalties that depend on the number of terms m in an m -term greedy approximation \hat{f}_m to f^* . We take \hat{f}_m to be an m term fit from an LPGP algorithm and choose \hat{m} among all $m \in \mathcal{M}$ (i.e. $\mathcal{M} = \{1, \dots, n\}$) to minimize

$$\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{f}_m(X_i))^2 + pen_n(\hat{f}_m, m)/n.$$

This approach enables the use of a data-based stopping criterion for the greedy algorithm. For more details on these adaptive methods, we refer the reader to [15]. The resolvability risk bound allows also for interpolation rates between L_2 and $L_{1,\mathcal{H}}$ refining the results of [8] and in accordance with the best balance between error of approximation and penalty.

The target f^* is not necessarily in \mathcal{F} . To each f in \mathcal{F} , there corresponds a function ρ , which assigns to (X, Y) the relative loss

$$\rho(X, Y) = \rho_f(X, Y) = (Y - f(X))^2 - (Y - f^*(X))^2.$$

Let \underline{X}' be an independent copy of the training data \underline{X} used for testing the efficacy of a fit \hat{f} based on $\underline{X}, \underline{Y}$. The relative empirical loss with respect to the training data is denoted by $P_n(f||f^*) = \frac{1}{n} \sum_{i=1}^n \rho(X_i, Y_i)$ and that with respect to the independent copy is $P'_n(f||f^*) = \frac{1}{n} \sum_{i=1}^n \rho(X'_i, Y_i)$. We define the empirical squared error by on the training and test data by $D_n(f, \tilde{f}) = \frac{1}{n} \sum_{i=1}^n (f(X_i) - \tilde{f}(X_i))^2$ and $D'_n(f, \tilde{f}) = \frac{1}{n} \sum_{i=1}^n (f(X'_i) - \tilde{f}(X'_i))^2$ for all f, \tilde{f} in \mathcal{F} . Using the relationship $Y = f^*(X) + \epsilon$, we note that $\rho(X, Y)$ can

also be written as $(f(X) - f^*(X))^2 - 2\epsilon(f(X) - f^*(X)) = g^2(X) - 2\epsilon g(X)$, where $g(x) = f(x) - f^*(x)$. Hence we have the relationship $P_n(f||f^*) = D_n(f, f^*) - \frac{2}{n} \sum_{i=1}^n \epsilon_i g(X_i)$.

The relative empirical loss $P'_n(f||f^*)$ is an unbiased estimate of the risk $\mathbb{E}\|\hat{f} - f^*\|^2$. Since ϵ_i has mean zero conditioned on X_i , the mean of $P'_n(f||f^*)$ with respect to \underline{X}' and \underline{Y} is $\mathbb{E}\|\hat{f} - f^*\|^2$. Such a quantity captures how well the fit \hat{f} based on the training data generalizes to a new set of observations. The goal is to control the empirical discrepancy $P'_n(f||f^*) - cP_n(f||f^*)$ between the loss on the future data and the loss on the training data for a constant $\tau > 1$. Toward this end, we seek a positive quantity $\text{pen}_n(f)$ to satisfy

$$\mathbb{E} \sup_{f \in \mathcal{F}} \{P'_n(f||f^*) - \tau P_n(f||f^*) - \tau \text{pen}_n(f)/n\} \leq 0,$$

Once such an inequality holds, the data-based choice \hat{f} in \mathcal{F} yields

$$\mathbb{E} P'_n(\hat{f}||f^*) \leq \tau \mathbb{E}[P_n(\hat{f}||f^*) + \text{pen}_n(f)/n].$$

If \hat{f} satisfies

$$(3.1) \quad \frac{1}{n} \sum_{i=1}^n (Y_i - \hat{f}(X_i))^2 + \frac{\text{pen}_n(\hat{f})}{n} \leq \inf_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n (Y_i - f(X_i))^2 + \frac{\text{pen}_n(f)}{n} + A_f \right\},$$

for some positive quantity A_f that decays to zero as the sample size grows, we see that

$$\mathbb{E} P'_n(\hat{f}||f^*) \leq \tau \inf_{f \in \mathcal{F}} \mathbb{E}[P_n(f||f^*) + \text{pen}_n(f)/n + A_f].$$

Using $\mathbb{E} P'_n(\hat{f}||f^*) = \mathbb{E}\|\hat{f} - f^*\|^2$ and $\mathbb{E} P_n(f||f^*) = \|f - f^*\|^2$, the above expression is seen to be

$$\mathbb{E}\|\hat{f} - f^*\|^2 \leq \tau \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E} \text{pen}_n(f)/n + \mathbb{E} A_f\}.$$

For the purposes of proving results in the case when \mathcal{F} is uncountable, it is useful to consider complexities $L_n(\tilde{f})$ for \tilde{f} in a countable subset $\tilde{\mathcal{F}}$ of \mathcal{F} satisfying $\sum_{\tilde{f} \in \tilde{\mathcal{F}}} e^{-\gamma_n L_n(\tilde{f})} \leq 1$ for some $\gamma_n > 0$ and such that

$$(3.2) \quad \begin{aligned} & \sup_{f \in \mathcal{F}} \{P'_n(f||f^*) - \tau P_n(f||f^*) - \tau \text{pen}_n(f)/n\} \\ & \leq \sup_{\tilde{f} \in \tilde{\mathcal{F}}} \{P'_n(\tilde{f}||f^*) - \tau P_n(\tilde{f}||f^*) - \tau \gamma_n L_n(\tilde{f})/n\}, \end{aligned}$$

with

$$\mathbb{E} \sup_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ P'_n(\tilde{f} \| f^*) - \tau P_n(\tilde{f} \| f^*) - \tau \gamma_n L_n(\tilde{f}) / n \right\} \leq 0.$$

The condition in (3.2) is equivalent to requiring that

$$\sup_{f \in \mathcal{F}} \inf_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ \Delta_n(f, \tilde{f}) + \gamma_n L_n(\tilde{f}) - \text{pen}_n(f) \right\} \leq 0,$$

where

$$\Delta_n(f, \tilde{f}) = n[P_n(\tilde{f} \| f^*) - P_n(f \| f^*)] - (n/\tau)[P'_n(\tilde{f} \| f^*) - P'_n(f \| f^*)].$$

If we truncate the penalized least squares estimator \hat{f} at a certain level B_n , for $\mathbb{E} \|T\hat{f} - f^*\|^2$ to maintain the resolvability bound $\tau \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E} \text{pen}_n(f) / n + \mathbb{E} A_f\}$, we require that

$$\sup_{f \in \mathcal{F}} \inf_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ \Delta_n(f, \tilde{f}) + \gamma_n L_n(\tilde{f}) - \text{pen}_n(f) \right\} \leq 0,$$

where

$$\Delta_n(f, \tilde{f}) = n[P_n(T\tilde{f} \| f^*) - P_n(f \| f^*)] - (n/\tau)[P'_n(T\tilde{f} \| f^*) - P'_n(Tf \| f^*)].$$

Rather than working with the relative empirical loss $P'_n(Tf \| f^*)$, we prefer to work with $D'_n(Tf, f^*)$. These two quantities are related to each other, provided $\frac{1}{n} \sum_{i=1}^n \epsilon_i g(X'_i)$ is small and they are exactly equal in the no noise case. Hence we would like to determine penalties that ensure

$$\mathbb{E} \sup_{f \in \mathcal{F}} \left\{ D'_n(Tf, f^*) - \tau P_n(f \| f^*) - \tau \text{pen}_n(f) / n \right\} \leq 0.$$

Suppose we require that

$$\begin{aligned} \mathbb{E} \sup_{f \in \mathcal{F}} \left\{ \tau_1^{-1} D'_n(Tf, f^*) - \tau \tau_2 P_n(f \| f^*) - \right. \\ \left. \tau \tau_2 \text{pen}_n(f) / n \right\} \leq 0, \end{aligned}$$

for some $\tau_1, \tau_2 \geq 1$. This further inflates the resulting risk bound by $(1 + \tau_1)(1 + \tau_2)$ so that the factor τ is replaced with $\tau \tau_1 \tau_2$ in (3). However, it enables us to create countable covers $\tilde{\mathcal{F}}$ with smaller errors in approximating functions from \mathcal{F} . To see this, suppose the countable cover $\tilde{\mathcal{F}}$ satisfies

$$\begin{aligned} \sup_{f \in \mathcal{F}} \left\{ \tau_1^{-1} D'_n(Tf, f^*) - \tau \tau_2 P_n(f \| f^*) - \tau \tau_2 \text{pen}_n(f) / n \right\} \\ \leq \sup_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ D'_n(T\tilde{f}, f^*) - \tau P_n(T\tilde{f} \| f^*) - \tau \gamma_n L_n(\tilde{f}) / n \right\}, \end{aligned}$$

or equivalently that

$$\sup_{f \in \mathcal{F}} \inf_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ \Delta_n(f, \tilde{f}) + \frac{\gamma_n}{\tau_2} L_n(\tilde{f}) - \text{pen}_n(f) \right\} \leq 0,$$

where

$$\begin{aligned} \Delta_n(f, \tilde{f}) &= n[\tau_2^{-1} P_n(T\tilde{f}||f^*) - P_n(f||f^*)] + \\ &\quad n\tau^{-1}\tau_2^{-1}[\tau_1^{-1} D'_n(Tf, f^*) - D'_n(T\tilde{f}, f^*)]. \end{aligned}$$

There are two cases to consider for bounding $\Delta_n(f, \tilde{f})$. In the noise case, we set $\tau_2 = 1$ and $\tau_1 = 1/\tau + 1$. Using the inequality, $\tau^{-1}a^2 - b^2 \leq \frac{1}{\tau-1}(b-a)^2$ that can be derived from $(a/\sqrt{\tau} - b\sqrt{\tau})^2 \geq 0$, we can upper bound the difference $\tau_1^{-1} D'_n(Tf, f^*) - D'_n(T\tilde{f}, f^*)$ by

$$(\tau_1 - 1)^{-1} D'_n(Tf, T\tilde{f}).$$

This quantity does not involve f^* , which is desirable for the proceeding analysis. Hence $\Delta_n(f, \tilde{f})$ is not greater than

$$n[P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})].$$

and thus we seek a penalty $\text{pen}_n(f)$ that is at least

$$\gamma_n L_n(\tilde{f}) + n[P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})].$$

An estimator \hat{f} satisfying (3.1) with penalty $\text{pen}_n(f)$ that is at least

$$\gamma_n L_n(\tilde{f}) + n[P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})].$$

satisfies the risk bound

$$\mathbb{E}\|T\hat{f} - f^*\|^2 \leq (\tau + 1) \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n + \mathbb{E}A_f\}.$$

In the no noise case, we set both τ_1 and τ_2 to be strictly greater than one. We can bound the difference

$$\begin{aligned} \tau_2^{-1} P_n(T\tilde{f}||f^*) - P_n(f||f^*) &= \tau_2^{-1} D_n(T\tilde{f}, f^*) - D_n(f, f^*) \\ &\leq \tau_2^{-1} D_n(\tilde{f}, f^*) - D_n(f, f^*) \end{aligned}$$

by $(\tau_2 - 1)^{-1} D_n(\tilde{f}, f)$ so that $\Delta_n(f, \tilde{f})$ has the bound

$$n(\tau_2 - 1)^{-1} D_n(\tilde{f}, f) + n\tau^{-1}\tau_2^{-1}(\tau_1 - 1)^{-1} D'_n(f, \tilde{f}).$$

If we set $\tau_2 = 2$ and $\tau_1 = 1/\tau + 1$, we find that $\Delta_n(f, \tilde{f})$ is no less than

$$n[D_n(\tilde{f}, f) + D'_n(f, \tilde{f})].$$

An estimator \hat{f} satisfying (3.1) with penalty $\text{pen}_n(f)$ that is at least

$$\gamma_n L_n(\tilde{f}) + n[D_n(\tilde{f}, f) + D'_n(f, \tilde{f})].$$

satisfies the risk bound

$$\mathbb{E}\|T\hat{f} - f^*\|^2 \leq 2(\tau + 1) \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n + \mathbb{E}A_f\}.$$

By bounding the distortion in this way, we eliminate some error in approximating f by \tilde{f} that arises from analyzing $P_n(T\tilde{f}||f^*) - P_n(f||f^*)$ and $D_n(T\tilde{f}, f^*) - D_n(Tf, f^*)$.

THEOREM 3. *Suppose $\tilde{\mathcal{F}}$ is a countable collection of functions that satisfies*

$$\mathbb{E} \sup_{\tilde{f} \in \tilde{\mathcal{F}}} \left\{ D'_n(T\tilde{f}, f^*) - \tau P_n(\tilde{f}||f^*) - \tau \gamma_n L_n(\tilde{f}) \right\} \leq 0.$$

If $\text{pen}_n(f)$ is at least

$$\gamma_n L_n(\tilde{f}) + n[P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})].$$

or

$$\gamma_n L_n(\tilde{f}) + n[D_n(\tilde{f}, f) + D'_n(f, \tilde{f})]$$

corresponding to the noise or no noise setting, then the truncated estimator $T\hat{f}$ with \hat{f} satisfying (3.1) has the resolvability bound

$$\mathbb{E}\|T\hat{f} - f^*\|^2 \leq 2(\tau + 1) \inf_{f \in \mathcal{F}} \{\|f - f^*\|^2 + \mathbb{E}\text{pen}_n(f)/n + \mathbb{E}A_f\}.$$

Recall that g is equal to $f - f^*$. In this way, there is a one to one correspondence between f and g . To simplify notation, we sometimes write $D_n(f, f^*)$ as $D_n(g)$ and $D'_n(f, f^*)$ as $D'_n(g)$. Moreover, assume an analogous notation holds for the relative loss functions $P_n(f||f^*)$ and $P'(f||f^*)$ and complexities $L_n(f)$.

THEOREM 4. *If \mathcal{F} is a countable collection of functions bounded in magnitude by B_n and $L_n(f)$ satisfies the Kraft inequality $\sum_{f \in \mathcal{F}} e^{-L_n(f)} \leq 1$, then*

$$\mathbb{E} \sup_{f \in \mathcal{F}} \left\{ D'_n(f||f^*) - \tau P_n(f||f^*) - \tau \gamma_n L_n(f)/n \right\} \leq 0,$$

where $\tau = (1 + \delta_1)(1 + \delta_2)$ and $\gamma_n = (2\tau)^{-1}(1 + \delta_1/2)(1 + 2/\delta_1)(B + B_n)^2 + 2(1 + 1/\delta_2)\sigma^2 + 2(B + B_n)\eta$.

PROOF. Let $s^2(g)$ be as in Lemma 1. Since g^2 is non-negative, $s^2(g) \leq D'_n(g^2) + D_n(g^2)$. Moreover, since $|f| \leq B_n$ and $|f^*| \leq B$, it follows that $s^2(g) \leq (B + B_n)^2(D'_n(g) + D_n(g))$. Let $\gamma_1 = A_1(B + B_n)^2/2$ with A_1 to be specified later. By Lemma 1, we have

$$(3.3) \quad \mathbb{E} \sup_{g \in \mathcal{G}} \left\{ (1 - 1/A_1)D'_n(g) - (1 + 1/A_1)D_n(g) - \frac{\gamma_1}{n}L(g) \right\}$$

$$(3.4) \quad \leq \mathbb{E} \sup_{g \in \mathcal{G}} \left\{ D'_n(g) - D_n(g) - \frac{\gamma_1}{n}L(g) - \frac{1}{2\gamma_1}s^2(g) \right\} \leq 0$$

By Lemma 2, we also know that

$$(3.5) \quad \mathbb{E} \sup_{g \in \mathcal{G}} \left\{ \frac{1}{n} \sum_{i=1}^n \epsilon_i g(X_i) - \frac{\gamma_2}{n}L(g) - \frac{1}{A_2 n}D_n(g) \right\} \leq 0,$$

where $\gamma_2 = A_2\sigma^2/2 + (B + B_n)\eta$. Adding the expression in (3.3) to $2a > 0$ times the expression in (3.5) and collecting terms, we find that $1 + 1/A_1 + 2a/A_2$ should be equal to a in order for $D_n(g)$ and $\frac{1}{n} \sum_{i=1}^n \epsilon_i g(X_i)$ to be added together to produce $P_n(g)$. Thus we find that

$$\mathbb{E} \sup_{g \in \mathcal{G}} \left\{ (1 - 1/A_1)D'_n(g) - a(P_n(g) + \frac{\gamma_n}{n}L(g)) \right\} \leq 0,$$

where $\gamma_n = \gamma_1/a + 2\gamma_2$. Choosing $A_1 = 1 + 2/\delta_1$, $A_2 = 2(1 + 1/\delta_2)$, and $\tau = (1 + \delta_1)(1 + \delta_2)$, we find that $a = \tau(1 - 1/A_1)$. Dividing the resulting expression by $1 - 1/A_1$ produces

$$\mathbb{E} \sup_{g \in \mathcal{G}} \left\{ D'_n(g) - \tau P_n(g) - \tau \gamma_n L(g)/n \right\} \leq 0.$$

□

In general, the penalty should not depend on the unknown test data \underline{X}' . However if one seeks to describe the error of a fit \hat{f} trained with the data $(\underline{X}, \underline{Y})$ at new data points \underline{X}' , a penalty that depends on \underline{X}' is natural and fits in with the standard trans-inductive setting in machine learning [13]. Since we have been assuming that the input design X is contained in a cube with side length at most one, the dependence on \underline{X}' of the penalties as shown in the following lemmata can be ignored.

When we speak of empirical L^2 covers of \mathcal{H} , we mean with respect to the empirical measure of $\underline{X} \cup \underline{X}'$ on both the training and test data. That is, empirical L^2 covers of \mathcal{H} are with respect to the squared norm $[D(h, \tilde{h}) + D'(h, \tilde{h})]/2$.

THEOREM 5. *Let $f = \sum_h \beta_h h$. Let $\tilde{\mathcal{H}}_1$ be an empirical L^2 ϵ_1 -net for \mathcal{H} of cardinality M_1 . Let $\tilde{\mathcal{H}}_2$ be an empirical L^2 ϵ_2 -net for \mathcal{H} of cardinality M_2 . Suppose these empirical covers do not depend on the underlying data. There exists a subset $\tilde{\mathcal{F}}$ of \mathcal{F} with cardinality at most $\binom{M_2+M_1+m_0}{M_1+m_0}$ such that for $v \geq v_f$ and $\tilde{v} = v(1 + M_1/m_0)$*

$$P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f}) \leq \frac{2\tilde{v}^2\epsilon_1^2}{m_0} + \frac{\tilde{v}^2 M_1}{2m_0^2} + 8B_n\tilde{v}\epsilon_2 + \frac{T_n}{n},$$

for some \tilde{f} in $\tilde{\mathcal{F}}$. Alternatively, there exists a subset $\tilde{\mathcal{F}}$ of \mathcal{F} with cardinality at most $\binom{M_2+m_0}{m_0}$ such that

$$P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f}) \leq \frac{2vv_f}{m_0} + 8B_nv\epsilon_2 + \frac{T_n}{n},$$

and in the case of no noise with $B_n \geq B$

$$D_n(\tilde{f}, f) + D'_n(f, \tilde{f}) \leq \frac{4vv_f}{m_0} + 4v^2\epsilon_2^2$$

for some \tilde{f} in $\tilde{\mathcal{F}}$.

PROOF. The proof is an immediate consequence of Lemma 4. \square

According to Theorems 3 and 4, a valid penalty is at least

$$\gamma_n L_n(\tilde{f}) + n[P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})],$$

where \tilde{f} belongs to a countable set $\tilde{\mathcal{F}}$ satisfying $\sum_{\tilde{f} \in \tilde{\mathcal{F}}} e^{-L_n(\tilde{f})} \leq 1$. The constant γ_n is as prescribed in Theorem 4. By Theorem 5, there is a set $\tilde{\mathcal{F}}$ with cardinality at most $\binom{M_2+M_1+m_0}{M_1+m_0}$ such that for all f with $v_f \leq v$, there is a \tilde{f} in $\tilde{\mathcal{F}}$ such that $P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})$ is bounded by

$$\frac{2\tilde{v}^2\epsilon_1^2}{m_0} + \frac{\tilde{v}^2 M_1}{2m_0^2} + 8B_n\tilde{v}\epsilon_2 + \frac{T_n}{n}.$$

Using the fact that the logarithm of $\binom{M_2+M_1+m_0}{M_1+m_0}$ is bounded by $(M_1 + m_0) \log(e(M_2/M_1 + 1))$, a valid penalty divided by sample size is at least

$$(3.6) \quad \frac{\gamma_n}{n} (M_1 + m_0) \log(e(M_2/M_1 + 1)) + \frac{2\tilde{v}^2\epsilon_1^2}{m_0} + \frac{\tilde{v}^2 M_1}{2m_0^2} + 8B_n\tilde{v}\epsilon_2 + \frac{T_n}{n}.$$

Alternatively, there is a set $\tilde{\mathcal{F}}$ with cardinality at most $\binom{M_2+m_0}{m_0}$ such that for all f with $v_f \leq v$, there is a \tilde{f} in $\tilde{\mathcal{F}}$ such that $P_n(T\tilde{f}||f^*) - P_n(f||f^*) + D'_n(Tf, T\tilde{f})$ is bounded by

$$\frac{2vv_f}{m_0} + 8B_nv\epsilon_2 + \frac{T_n}{n}$$

and hence a valid penalty divided by sample size is at least

$$(3.7) \quad \frac{\gamma_n m_0 \log M_2}{n} + \frac{2vv_f}{m_0} + 8B_nv\epsilon_2 + \frac{T_n}{n}.$$

In the no noise case, a valid penalty divided by sample size is at least

$$(3.8) \quad \frac{\gamma_n m_0 \log M_2}{n} + \frac{4vv_f}{m_0} + 4v^2\epsilon_2^2.$$

We now discuss how m_0 , ϵ_1 , and ϵ_2 should be chosen to produce penalties that yield optimal risk properties for $T\hat{f}$.

4. Risk bounds in high dimensions.

4.1. *Noise case.* By Lemma 6, an empirical L^2 ϵ_2 -cover of \mathcal{H} has cardinality less than $\binom{2d+\lceil(\Lambda/\epsilon_2)^2\rceil}{\lceil(\Lambda/\epsilon_2)^2\rceil}$. The logarithm of $\binom{2d+\lceil(\Lambda/\epsilon_2)^2\rceil}{\lceil(\Lambda/\epsilon_2)^2\rceil}$ is bounded by $4(\Lambda/\epsilon_2)^2 \log(d+1)$.

Continuing from the expression (3.7), we find that $\text{pen}_n(f)/n$ is at least

$$\frac{4\gamma_n m_0 (\Lambda/\epsilon_2)^2 \log(d+1)}{n} + \frac{2vv_f}{m_0} + 8B_nv\epsilon_2 + \frac{T_n}{n}.$$

Choosing m_0 to be the ceiling of $\left(\frac{vv_f n \epsilon_2^2}{2\gamma_n \Lambda^2 \log(d+1)}\right)^{1/2}$, we see that $\text{pen}_n(f)/n$ must be at least

$$\frac{8\gamma_n \Lambda^2 \log(d+1)}{n\epsilon_2^2} + 8 \left(\frac{vv_f \gamma_n \Lambda^2 \log(d+1)}{n\epsilon_2^2} \right)^{1/2} + 8B_nv\epsilon_2 + \frac{T_n}{n}.$$

Finally, we set $v = v_f$ and $\epsilon_2 = \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{nB_n^2}\right)^{1/4}$ so that $\text{pen}_n(f)/n$ must be at least

$$16v_f \left(\frac{\gamma_n B_n^2 \Lambda^2 \log(d+1)}{n} \right)^{1/4} + 8 \left(\frac{\gamma_n B_n^2 \Lambda^2 \log(d+1)}{n} \right)^{1/2} + \frac{T_n}{n}.$$

We see that the main term in the penalty divided by sample size is

$$16v_f \left(\frac{\gamma_n B_n^2 \Lambda^2 \log(d+1)}{n} \right)^{1/4}.$$

4.2. *No noise case.* Continuing from the expression (3.8), we find that $\text{pen}_n(f)/n$ is at least

$$\frac{4\gamma_n m_0 (\Lambda/\epsilon_2)^2 \log(d+1)}{n} + \frac{4vv_f}{m_0} + 4v^2 \epsilon_2^2.$$

Choosing m_0 to be the ceiling of $\left(\frac{vv_f n \epsilon_2^2 (\tau+1)^2}{\gamma_n \Lambda^2 \log(d+1) (\tau+2)} \right)^{1/2}$, we see that $\text{pen}_n(f)/n$ must be at least

$$\frac{4\gamma_n \Lambda^2 \log(d+1)}{n \epsilon_2^2} + 8 \left(\frac{vv_f \gamma_n \Lambda^2 \log(d+1)}{n \epsilon_2^2} \right)^{1/2} + 4v^2 \epsilon_2^2.$$

Finally, we set $v = v_f$ and $\epsilon_2 = \left(\frac{4(\tau+2)/(\tau+1)^2 \gamma_n \Lambda^2 \log(d+1)}{nv^2} \right)^{1/6}$ so that $\text{pen}_n(f)/n$ must be at least

$$16v_f^{4/3} \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{n} \right)^{1/3} + 4(v_f^{4/3} + 1) \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{n} \right)^{2/3},$$

where we used the fact that $v^{2/3} \leq v^{4/3} + 1$. We see that the main term in the penalty divided by sample size is

$$16v_f^{4/3} \left(\frac{\gamma_n \Lambda^2 \log(d+1)}{n} \right)^{1/3}.$$

4.3. *Combining the noise and no noise cases.* As the previous sections show, there are differences in the risk bounds depending on the nature of the noise. These bounds only incorporate information about the presence of noise, without regard to the degree of variability. Ideally one would prefer to have a bound that interpolates between these two situations and recovers the no noise case exactly when the variability is zero. We can establish the validity of a penalty divided by sample size that is at least a constant multiple of

$$(4.1) \quad \left(\frac{v_f^4 \Lambda^2 \gamma_n \log(d+1)}{n} \right)^{1/3} + \sqrt{\sigma} \left(\frac{v_f^4 \Lambda^2 \gamma_n \log(d+1)}{n} \right)^{1/4}.$$

plus negligible terms that do not depend on the candidate fit. Note that this penalty is a convex function of v_f and hence the setting of Theorem 1 applies. For a proof of the above form of the penalty, we refer the reader to the supplement [Supplement A](#).

5. Risk bounds with improved exponents for moderate dimensions. Continuing from the expression (3.6), we find that $\text{pen}_n(f)/n$ is at least

$$\frac{\gamma_n}{n}(M_1 + m_0) \log(e(M_2/M_1 + 1)) + \frac{2\tilde{v}^2\epsilon_1^2}{m_0} + \frac{\tilde{v}^2 M_1}{2m_0^2} + 8B_n\tilde{v}\epsilon_2 + \frac{T_n}{n}.$$

Note that we can bound B_n^2 by γ_n by choosing δ_1 and δ_2 appropriately. For the precise definition of γ_n , see Theorem 4. The strategy for optimization is to first consider the terms

$$(5.1) \quad \frac{\gamma_n}{n}m_0 \log(e(M_2/M_1 + 1)) + \frac{2\tilde{v}^2\epsilon_1^2}{m_0} + 8\sqrt{\gamma_n}\tilde{v}\epsilon_2.$$

After m_0 , M_1 , and M_2 have been selected, we then check that

$$(5.2) \quad \frac{\gamma_n}{n}M_1 \log(e(M_2/M_1 + 1)) + \frac{\tilde{v}^2 M_1}{2m_0^2}$$

is relatively negligible. Choosing m_0 to be the ceiling of $\left(\frac{2\tilde{v}^2 n \epsilon_1^2}{\gamma_n \log(e(M_2/M_1 + 1))}\right)^{1/2}$, we see that (5.1) is at most

$$\frac{\gamma_n}{n} \log(e(M_2/M_1 + 1)) + 4 \left(\frac{\tilde{v}^2 \gamma_n \epsilon_1^2 \log(e(M_2/M_1 + 1))}{n} \right)^{1/2} + 8\sqrt{\gamma_n}\tilde{v}\epsilon_2.$$

Note that an empirical L^2 ϵ -cover of \mathcal{H} has cardinality between $(\Lambda/\epsilon)^d$ and $(2\Lambda/\epsilon + 1)^d \leq (3\Lambda/\epsilon)^d$ whenever $\epsilon \leq \Lambda$. Thus $M_2/M_1 \leq (3\epsilon_1/\epsilon_2)^d$ whenever $\epsilon_2 \leq \Lambda$ and hence

$$\log(e(M_2/M_1 + 1)) \leq 1 + (d/2) \log(9\epsilon_1^2/\epsilon_2^2 + 1) \leq d \log(9\epsilon_1^2/\epsilon_2^2 + 1),$$

whenever $\epsilon_1^2 \geq \epsilon_2^2(e-1)/9$. These inequalities imply that (5.1) is at most

$$\frac{d\gamma_n \log(9\epsilon_1^2/\epsilon_2^2 + 1)}{n} + 4 \left(\frac{\tilde{v}^2 \epsilon_1^2 d \gamma_n \log(9\epsilon_1^2/\epsilon_2^2 + 1)}{n} \right)^{1/2} + 8\sqrt{\gamma_n}\tilde{v}\epsilon_2.$$

Next, set

$$\epsilon_2^2 = \frac{9d\epsilon_1^2}{n}.$$

This means that the assumption $\epsilon_1^2 \geq \epsilon_2^2(e-1)/9$ is valid provided $d \leq n/(e-1)$. Thus (5.1) is at most

$$\frac{d\gamma_n \log(n/d+1)}{n} + 20\epsilon_1 \tilde{v} \sqrt{\frac{d\gamma_n \log(n/d+1)}{n}}.$$

Next, we add in the terms from (5.2). The selections of m_0 and ϵ_1 make (5.2) at most

$$\frac{M_1 d\gamma_n \log(n/d+1)}{n} + \frac{M_1 d\gamma_n \log(n/d+1)}{n\epsilon_1^2}$$

Since $M_1 \leq (3\Lambda/\epsilon_1)^d$ whenever $\epsilon_1 \leq \Lambda$, we find that (5.2) is at most

$$\frac{(3\Lambda)^d d\gamma_n \log(n/d+1)}{n\epsilon_1^d} + \frac{(3\Lambda)^d d\gamma_n \log(n/d+1)}{n\epsilon_1^{d+2}}$$

Let $\epsilon_1 = 3\Lambda \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2(d+3)}$. Choosing $\tilde{v} = v_f$, we see that a valid penalty divided by sample size is at least

$$60v_f \Lambda \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+1/2(d+3)} + \frac{1}{\Lambda^2} \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+1/2(d+3)} \\ + \left(\frac{d\gamma_n \log(n/d+1)}{n} \right)^{1/2+3/2(d+3)} + \frac{d\gamma_n \log(n/d+1)}{n} + \frac{T_n}{n}.$$

Note that for the form of the above penalty to be valid, we need $\frac{d\gamma_n \log(n/d)}{n}$ to be small enough to ensure that ϵ_1 and ϵ_2 are both less than Λ .

REMARK. *It is possible to remove the term $\frac{v^2 M_1}{4m_0^2}$ from the bound in Lemma 5 at the expense of a larger set $\tilde{\mathcal{F}}$. For example, let $(1/m_0) \sum_{k=1}^m b_k h_k$, $\|b\| \geq v_f$ be as in the first part of Lemma 4. Note that the vector $(b_1, \dots, b_m)' / \|b\|_1$ belongs to the ℓ_1 unit ball. This space has an ϵ -covering number of order $(1/\epsilon)^{m_0+M_1}$ and thus the collection of representors $\tilde{\mathcal{F}}$ is of order*

$$(1/\epsilon_2)^{m_0+M_1} \binom{M_2 + M_1 + m_0}{M_1 + m_0}.$$

This new cover can be used to yield even tighter bounds than those presented in Table 1. However, the analysis is more technical and we omit it here. We

will however say that slightly improved rates of order $v_{f^*} \left(\frac{d}{n} \log \frac{n}{d}\right)^{1/2+1/(2(d+1))}$ are possible.

6. Proofs of the lemmata. An important aspect of the above covers $\tilde{\mathcal{F}}$ is that they only depend on the data $(\underline{X}, \underline{X}')$ through $\|\underline{X}\|_\infty^2 + \|\underline{X}'\|_\infty^2$, where $\|\underline{X}\|_\infty^2 = \frac{1}{n} \sum_{i=1}^n \|X_i\|_\infty^2$. Since the coordinates of \underline{X} and \underline{X}' are restricted to belong to $[-1, 1]^d$, the penalties and quantities satisfying Kraft's inequality do not depend on \underline{X} and \underline{X}' . This is an important implication for the following empirical process theory. On the other hand, using the fact that $\|\underline{X}\|_\infty^2 + \|\underline{X}'\|_\infty^2$ is symmetric in the coordinates of \underline{X} and \underline{X}' and has a mean that is at most logarithmic in d , the following bounds can be adapted to handle covers $\tilde{\mathcal{F}}$ that depend on the training and test data without imposing sup-norm controls.

LEMMA 1. Let $(\underline{X}, \underline{X}') = (X_1, \dots, X_n, X'_1, \dots, X'_n)$, where \underline{X}' is an independent copy of the data \underline{X} and where (X_1, \dots, X_n) are component-wise independent but not necessarily identically distributed. A countable function class \mathcal{G} and complexities $L(g)$ satisfying $\sum_{g \in \mathcal{G}} e^{-L(g)} \leq 1$ are given. Then for arbitrary positive γ ,

$$(6.1) \quad \mathbb{E} \sup_{g \in \mathcal{G}} \left\{ D'_n(g) - D_n(g) - \frac{\gamma}{n} L(g) - \frac{1}{2\gamma} s^2(g) \right\} \leq 0,$$

where $s^2(g) = \frac{1}{n} \sum_{i=1}^n (g^2(X_i) - g^2(X'_i))^2$.

PROOF. Let $\underline{Z} = (Z_1, \dots, Z_n)$ be a sequence of independent centered Bernoulli random variables with success probability $1/2$. Since X_i and X'_i are identically distributed, $g^2(X_i) - g^2(X'_i)$ is a symmetric random variable and hence sign changes do not affect the expectation in (6.1). Thus the right hand side of the inequality in (6.1) is equal to

$$\mathbb{E}_{\underline{Z}, \underline{X}, \underline{X}'} \sup_{g \in \mathcal{G}} \left\{ \frac{1}{n} \sum_{i=1}^n Z_i (g^2(X_i) - g^2(X'_i)) - \frac{\gamma}{n} L(g) - \frac{1}{2\gamma} s^2(g) \right\}.$$

Using the identity $x = \lambda \log(x/\lambda)$ with $\lambda = \gamma/n$, conditioning on \underline{X} and \underline{X}' , and applying Jensen's inequality to move $\mathbb{E}_{\underline{Z}}$ inside the logarithm, we have

that

$$\begin{aligned} & \mathbb{E}_{\underline{Z}} \sup_{g \in \mathcal{G}} \left\{ \frac{1}{n} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) - \frac{\gamma}{n} L(g) - \frac{1}{2\gamma} s^2(g) \right\} \\ & \leq \frac{\gamma}{n} \log \mathbb{E}_{\underline{Z}} \sup_{g \in \mathcal{G}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) - L(g) - \frac{n}{2\gamma^2} s^2(g) \right\}. \end{aligned}$$

Replacing the supremum with the sum and using the linearity of expectation, the above expression is not more than

$$\begin{aligned} & \frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} \mathbb{E}_{\underline{Z}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) - L(g) - \frac{n}{2\gamma^2} s^2(g) \right\} \\ & = \frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} \exp \left\{ -L(g) - \frac{n}{2\gamma^2} s^2(g) \right\} \mathbb{E}_{\underline{Z}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) \right\}. \end{aligned}$$

Next, note that by the independence of Z_1, \dots, Z_n ,

$$\mathbb{E}_{\underline{Z}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) \right\} = \prod_{i=1}^n \mathbb{E}_{Z_i} \exp \left\{ \frac{1}{\gamma} Z_i(g^2(X_i) - g^2(X'_i)) \right\}.$$

Using the inequality $e^x + e^{-x} \leq 2e^{x^2/2}$, each $\mathbb{E}_{Z_i} \exp \left\{ \frac{1}{\gamma} Z_i(g^2(X_i) - g^2(X'_i)) \right\}$ is not more than $\exp \left\{ \frac{1}{2\gamma^2} (g^2(X_i) - g^2(X'_i))^2 \right\}$. Whence

$$\mathbb{E}_{\underline{Z}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n Z_i(g^2(X_i) - g^2(X'_i)) \right\} \leq \exp \left\{ \frac{n}{2\gamma^2} s^2(g) \right\}.$$

The claim follows from the fact that $\frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} e^{-L(g)} \leq 0$. \square

LEMMA 2. *Let $\underline{\epsilon} = (\epsilon_1, \dots, \epsilon_n)$ be conditionally independent random variables given $\{X_i\}_{i=1}^n$, with conditional mean zero, satisfying Bernstein's moment condition with parameter $\eta > 0$. A countable class \mathcal{G} and complexities $L(g)$ satisfying*

$$\sum_{g \in \mathcal{G}} e^{-L(g)} \leq 1$$

are given. Assume a bound K , such that $|g(x)| \leq K$ for all g in \mathcal{G} . Then

$$\mathbb{E} \sup_{g \in \mathcal{G}} \left\{ \frac{1}{n} \sum_{i=1}^n \epsilon_i g(X_i) - \frac{\gamma}{n} L(g) - \frac{1}{An} \sum_{i=1}^n g^2(X_i) \right\} \leq 0.$$

where A is an arbitrary constant and $\gamma = A\sigma^2/2 + Kh$.

PROOF. Using the identity $x = \lambda \log(x/\lambda)$ with $\lambda = \gamma/n$, conditioning on \underline{X} , and applying Jensen's inequality to move $\mathbb{E}_{\underline{\epsilon}}$ inside the logarithm, we have that

$$\begin{aligned} & \mathbb{E}_{\underline{\epsilon}|\underline{X}} \sup_{g \in \mathcal{G}} \left\{ \frac{1}{n} \sum_{i=1}^n \epsilon_i g(X_i) - \frac{\gamma}{n} L(g) - \frac{1}{An} \sum_{i=1}^n g^2(X_i) \right\} \\ & \leq \frac{\gamma}{n} \log \mathbb{E}_{\underline{\epsilon}|\underline{X}} \sup_{g \in \mathcal{G}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n \epsilon_i g(X_i) - L(g) - \frac{1}{\gamma A} \sum_{i=1}^n g^2(X_i) \right\}. \end{aligned}$$

Replacing the supremum with the sum and using the linearity of expectation, the above expression is not more than

$$\begin{aligned} & \frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} \mathbb{E}_{\underline{\epsilon}|\underline{X}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n \epsilon_i g(X_i) - L(g) - \frac{1}{\gamma A} \sum_{i=1}^n g^2(X_i) \right\} \\ & = \frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} \exp \left\{ -L(g) - \frac{1}{\gamma A} \sum_{i=1}^n g^2(X_i) \right\} \mathbb{E}_{\underline{\epsilon}|\underline{X}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n \epsilon_i g(X_i) \right\}. \end{aligned}$$

Next, note that by the independence of $\epsilon_1, \dots, \epsilon_n$ conditional on \underline{X} ,

$$\mathbb{E}_{\underline{\epsilon}|\underline{X}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n \epsilon_i g(X_i) \right\} = \prod_{i=1}^n \mathbb{E}_{\epsilon_i|X_i} \exp \left\{ \frac{1}{\gamma} \epsilon_i g(X_i) \right\}.$$

By Lemma 7, each $\mathbb{E}_{\epsilon_i|X_i} \exp \left\{ \frac{1}{\gamma} \epsilon_i g(X_i) \right\}$ is not more than $\exp \left\{ \frac{\sigma^2 g^2(X_i)}{2\gamma^2(1-\eta K/\gamma)} \right\}$. Whence

$$\begin{aligned} \mathbb{E}_{\underline{\epsilon}|\underline{X}} \exp \left\{ \frac{1}{\gamma} \sum_{i=1}^n \epsilon_i g(X_i) \right\} & \leq \exp \left\{ \frac{\sigma^2 \sum_{i=1}^n g^2(X_i)}{2\gamma^2(1-\eta K/\gamma)} \right\} \\ & = \exp \left\{ \frac{1}{\gamma A} \sum_{i=1}^n g^2(X_i) \right\}, \end{aligned}$$

where the last line follows from the definition of γ . The proof is finished after observing that $\frac{\gamma}{n} \log \sum_{g \in \mathcal{G}} e^{-L(g)} \leq 0$. \square

LEMMA 3. For $f = \sum_h \beta_h h$ and f_0 in \mathcal{F} , there is a choice of h_1, \dots, h_m in \mathcal{H} with $f_m = (v/m) \sum_{k=1}^m h_k$, $v \geq v_f$ such that

$$\|f_m - f_0\|^2 - \|f_0 - f\|^2 \leq \frac{v v_f}{m}.$$

Moreover, the same bound holds for any convex combination of $\|f_m - f_0\|^2 - \|f_0 - f\|^2$ and $\rho^2(f_m, f)$, where ρ is a possibly different Hilbert space norm.

PROOF. Let H be a random variable that equals hv with probability β_h/v and zero with probability $1 - v_f/v$. Let H_1, \dots, H_m be a random sample from the distribution defining H . Then $\overline{H} = \frac{1}{m} \sum_{j=1}^m H_j$ has mean f and furthermore the mean of $\|f_m - f_0\|^2 - \|f_0 - f\|^2$ is the mean is $\|f - \overline{H}\|^2$. This quantity is seen to be bounded by vv_f/m . As a consequence of the bound holding on average, there exists a realization of f_m of \overline{H} (having form $(v/m) \sum_{k=1}^m h_k$) such that $\|f_m - f_0\|^2 - \|f_0 - f\|^2$ is also bounded by Vv_f/m . \square

The next lemma is an extension of a technique used in [19] to improve the L^2 error of an m -term approximation of a function in $L_{1,\mathcal{H}}$. The idea is essentially stratified sampling with proportional allocation used in survey sampling as a means of variance reduction. In the following, we use the notation $\|\cdot\|$ to denote a generic Hilbert space norm.

LEMMA 4. *Let $\tilde{\mathcal{H}}$ be an L^2 ϵ_1 -net of \mathcal{H} with cardinality M_1 . For $f = \sum_h \beta_h h$ and f_0 in \mathcal{F} , there is a choice of h_1, \dots, h_m in \mathcal{H} with $f_m = (1/m_0) \sum_{k=1}^m b_k h_k$, $m \leq m_0 + M_1$ and $\|b\|_1 \geq v_f$ such that*

$$\|f_0 - f_m\|^2 - \|f_0 - f\|^2 \leq \frac{vv_f \epsilon_1^2}{m_0}.$$

Moreover, there is an equally weighted linear combination $f_m = (v/m_0) \sum_{k=1}^m h_k$, $v \geq v_f$, $m \leq m_0 + M_1$ such that

$$\|f_0 - f_m\|^2 - \|f_0 - f\|^2 \leq \frac{v^2 \epsilon_1^2 (1 + M_1/m_0)}{m_0} + \frac{v^2 M_1}{4m_0^2}.$$

The same bound holds for any convex combination of $\|f_m - f_0\|^2 - \|f_0 - f\|^2$ and $\rho^2(f_m, f)$, where ρ is a possibly different Hilbert space norm.

PROOF. Suppose the elements of $\tilde{\mathcal{H}}$ are $\tilde{h}_1, \dots, \tilde{h}_{M_1}$. Consider the M_1 sets

$$\tilde{\mathcal{H}}_j = \{h \in \mathcal{H} : \|h - \tilde{h}_j\|^2 \leq \epsilon_1^2\},$$

$j = 1, \dots, M_1$. By working instead with disjoint sets $\tilde{\mathcal{H}}_j \setminus \bigcup_{1 \leq i \leq j-1} \tilde{\mathcal{H}}_i$ ($\tilde{\mathcal{H}}_0 = \emptyset$) that are contained in $\tilde{\mathcal{H}}_j$ and whose union is \mathcal{H} , we may assume that the $\tilde{\mathcal{H}}_j$ form a partition of \mathcal{H} . Let $M = m_0 + M_1$ and $v_j = \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h$. To obtain the first conclusion, define a random variable H_j to equal hv_j with probability β_h/v_j for all $h \in \tilde{\mathcal{H}}_j$. let $H_{1,j}, \dots, H_{n_j,j}$ be a random sample

of size $N_j = \left\lceil \frac{v_j M}{V} \right\rceil$, where $V = \frac{vM}{m_0}$ and $v \geq v_f$, from the distribution defining H_j . Note that the N_j sum to at most M . Define $g_j = \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h h$ and $\bar{f} = \sum_{j=1}^{M_1} \frac{1}{N_j} \sum_{k=1}^{N_j} H_{k,j}$. Note that the mean of \bar{f} is f . This means the expectation of $\|f_0 - \bar{f}\|^2 - \|f_0 - f\|^2$ is the expectation of $\|f - \bar{f}\|^2$, which is equal to $\sum_{j=1}^{M_1} \mathbb{E} \|H_j - g_j\|^2 / N_j$. Now $\mathbb{E} \|H_j - g_j\|^2 / N_j$ is further bounded by

$$(V/M) \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h \inf_{h_j} \|h - h_j\|^2 \leq (V/M) \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h \|h - \tilde{h}_j\|^2 \leq \frac{v_j v \epsilon_1^2}{m_0}.$$

The above fact was established by noting that the mean of a real-valued random variable minimizes its average squared distance from any point h_j . Summing over $1 \leq j \leq M_1$ produces the claim. Since this bound holds on average, there exists a realization f_m of \bar{f} (having form $(1/m_0) \sum_{k=1}^m b_k h_k$ with $\|b\|_1 \geq v_f$) such that $\|f_0 - f_m\|^2 - \|f_0 - f\|^2$ is also bounded by $\frac{v v_f \epsilon_1^2}{m_0}$.

For the second conclusion, we proceed in a similar fashion. Suppose n_j is a random variable that equals $\left\lceil \frac{v_j M}{V} \right\rceil$ and $\left\lfloor \frac{v_j M}{V} \right\rfloor$ with respective probabilities chosen to make its average equal to $\frac{v_j M}{V}$. Furthermore, assume n_1, \dots, n_{M_1} are independent. Define $V_j = \frac{V}{M} n_j$. Since $V_j \leq v_j + \frac{V}{M}$, the V_j sum to at most V . Let H_j be a random variable that equals $h v_j$ with probability β_h / v_j for all $h \in \tilde{\mathcal{H}}_j$. For each j and conditional on n_j , let $H_{1,j}, \dots, H_{n_j,j}$ be a random sample of size $N_j = n_j + \mathbb{I}\{n_j = 0\}$ from the distribution defining H_j . Note that the N_j sum to at most M . Define $g_j = \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h h$ and $\bar{f} = \sum_{j=1}^{M_1} \frac{1}{N_j} \sum_{k=1}^{N_j} H_{k,j}$. Note that the conditional mean of \bar{H} given N_1, \dots, N_{M_1} is $g = \sum_{j=1}^{M_1} (V_j / v_j) g_j$ and hence the mean of \bar{f} is f . This means the expectation of $\|f_0 - \bar{f}\|^2 - \|f_0 - f\|^2$ is the expectation of $\|f - \bar{f}\|^2$, which is equal to $\sum_{j=1}^{M_1} \mathbb{E} \|H_j - (V_j / v_j) g_j\|^2 / N_j + \mathbb{E} \|f - g\|^2$ by the law of total variance. Now $\mathbb{E} \|H_j - (V_j / v_j) g_j\|^2 / N_j$ is further bounded by

$$(V/M)^2 (n_j / v_j) \sum_{h \in \tilde{\mathcal{H}}_j} \beta_h \inf_{h_j} \|h - h_j\|^2 \leq \frac{v^2 M \epsilon_1^2}{m_0^2}.$$

The above fact was established by noting that the mean of a real-valued random variable minimizes its average squared distance from any point h_j . Next, note that by the independence of the coordinates of v_1, \dots, v_{M_1} and

the fact that V_j has mean v_j ,

$$\mathbb{E}\|f - g\|^2 = \mathbb{E}\left\|\sum_{j=1}^{M_1} (v_j/v_j - 1)g_j\right\|^2 = (V/M)^2 \sum_{j=1}^{M_1} (\|g_j\|^2/v_j^2)\mathbb{V}(n_j).$$

Finally, observe that $\|g_j\|^2 \leq v_j^2$ and $\mathbb{V}(n_j) \leq 1/4$ (a random variable whose range is contained in an interval of length one has variance bounded by $1/4$). This shows that $\mathbb{E}\|f - g\|^2 \leq \frac{v^2 M_1}{4m_0^2}$. Since this bound holds on average, there exists a realization f_m of \bar{f} (having form $(v/m_0) \sum_{k=1}^m h_k$) such that $\|f_0 - f_m\|^2 - \|f_0 - f\|^2$ is also bounded by $\frac{v^2 \epsilon_1^2 (1 + M_1/m_0)}{m_0} + \frac{v^2 M_1}{4m_0^2}$. \square

LEMMA 5. *Let $\underline{y} = \{y_i\}_{i=1}^n$ be a sequence of real numbers and let $\underline{x} = \{x_i\}_{i=1}^n$ and $\underline{x}' = \{x'_i\}_{i=1}^n$ be sequences d -dimensional vectors. Let $\tilde{\mathcal{H}}_1$ be an empirical L^2 ϵ_1 -net for \mathcal{H} with cardinality M_1 and $\tilde{\mathcal{H}}_2$ be an empirical L^2 ϵ_2 -net for \mathcal{H} with cardinality M_2 . For $f = \sum_h \beta_h h$ in \mathcal{F} , there is a choice of $\tilde{h}_1, \dots, \tilde{h}_m$ in $\tilde{\mathcal{H}}_2$ with $\tilde{f}_m = (v/m_0) \sum_{k=1}^m \tilde{h}_k$, $v \geq v_f$, and $m \leq m_0 + M_1$, such that*

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n (y_i - T\tilde{f}_m(x_i))^2 - \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 + \\ & \quad \frac{1}{n} \sum_{i=1}^n (T\tilde{f}_m(x'_i))^2 - T f(x'_i))^2 \\ & \leq \frac{2v^2 \epsilon_1^2 (1 + M_1/m_0)}{m_0} + \frac{v^2 M_1}{2m_0^2} \\ & \quad + 8B_n v (1 + M_1/m_0) \epsilon_2 \\ & \quad + \frac{T_n}{n}. \end{aligned}$$

If $\tilde{\mathcal{F}}$ denotes the collection of functions of the form \tilde{f}_m , then $\tilde{\mathcal{F}}$ has cardinality at most $\binom{M_2 + M_1 + m_0}{M_1 + m_0}$.

Moreover, there is a choice of $\tilde{h}_1, \dots, \tilde{h}_{m_0}$ in $\tilde{\mathcal{H}}_2$ with $\tilde{f}_{m_0} = (v/m_0) \sum_{k=1}^{m_0} \tilde{h}_k$,

$v \geq v_f$ such that

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n (y_i - T\tilde{f}_{m_0}(x_i))^2 - \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 + \\
& \quad \frac{1}{n} \sum_{i=1}^n (T\tilde{f}_{m_0}(x'_i) - Tf(x'_i))^2 \\
(6.2) \quad & \leq \frac{2vv_f}{m_0} + 8B_n v \epsilon_2 + \frac{T_n}{n}.
\end{aligned}$$

and

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n (\tilde{f}_{m_0}(x_i) - f(x_i))^2 + \frac{1}{n} \sum_{i=1}^n (\tilde{f}_{m_0}(x'_i) - f(x'_i))^2 \\
(6.3) \quad & \leq \frac{4vv_f}{m_0} + 4v^2 \epsilon_2^2.
\end{aligned}$$

If $\tilde{\mathcal{F}}$ denotes the collection of functions of the form \tilde{f}_{m_0} , then $\tilde{\mathcal{F}}$ has cardinality at most $\binom{M_2+m_0}{m_0}$.

PROOF. We only prove the first claim of the lemma. Inequalities (6.2) and (6.3) follow from similar arguments and Lemma 3. Let $f_m = (v/m_0) \sum_{k=1}^m h_k$ be as in the second part of Lemma 4. Since $\tilde{\mathcal{H}}_2$ is an empirical L^2 ϵ_2 -net for \mathcal{H} , for each h_k there is an \tilde{h}_k in $\tilde{\mathcal{H}}_2$ such that

$$\frac{1}{2n} \sum_{i=1}^n |h_k(x_i) - \tilde{h}_k(x_i)|^2 + \frac{1}{2n} \sum_{i=1}^n |h_k(x'_i) - \tilde{h}_k(x'_i)|^2 \leq \epsilon_2^2.$$

Let $\tilde{f}_m = (v/m_0) \sum_{k=1}^m \tilde{h}_k$. By Lemma 8 (I) and (II),

$$\begin{aligned}
 (y - T\tilde{f}_m(x))^2 - (y - f(x))^2 &= [(y - f_m(x))^2 - (y - f(x))^2] + \\
 &\quad [(y - T\tilde{f}_m(x))^2 - (y - Tf_m(x))^2] + \\
 &\quad [(y - Tf_m(x))^2 - (y - f_m(x))^2] \\
 &\leq [(y - f_m(x))^2 - (y - f(x))^2] + \\
 &\quad 4B_n|f_m(x) - \tilde{f}_m(x)| + \\
 &\quad 4B_n(|y| - B_n)\mathbb{I}\{|y| > B_n\} + \\
 &\quad 2(|y| - B_n)^2\mathbb{I}\{|y| > B_n\} \\
 &= [(y - f_m(x))^2 - (y - f(x))^2] + \\
 &\quad 4B_n|f_m(x) - \tilde{f}_m(x)| + \\
 &\quad 2(|y|^2 - B_n^2)\mathbb{I}\{|y| > B_n\}.
 \end{aligned}$$

By Lemma 8 (III),

$$(T\tilde{f}_m(x') - Tf(x'))^2 \leq (f(x') - f_m(x'))^2 + 4B_n|\tilde{f}_m(x') - f_m(x')|.$$

Thus we find that $(y - T\tilde{f}_m(x))^2 - (y - f(x))^2 + (T\tilde{f}_m(x') - Tf(x'))^2$ is not greater than

$$\begin{aligned}
 &[(y - f_m(x))^2 - (y - f(x))^2] + (f(x') - f_m(x'))^2 + \\
 &4B_n[|f_m(x) - \tilde{f}_m(x)| + |\tilde{f}_m(x') - f_m(x')|] + 2(|y|^2 - B_n^2)\mathbb{I}\{|y| > B_n\}
 \end{aligned}$$

By the second conclusion in Lemma 4,

$$\begin{aligned}
 &\frac{1}{n} \sum_{i=1}^n (y_i - f_m(x_i))^2 - \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 + \\
 &\frac{1}{n} \sum_{i=1}^n (f_m(x'_i) - f(x'_i))^2 \leq \frac{2v^2\epsilon_1^2(1 + M_1/m_0)}{m_0} + \frac{v^2M_1}{2m_0^2}.
 \end{aligned}$$

By the concavity of the square root function,

$$\frac{1}{2n} \sum_{i=1}^n |h_k(x_i) - \tilde{h}_k(x_i)| + \frac{1}{2n} \sum_{i=1}^n |h_k(x'_i) - \tilde{h}_k(x'_i)|$$

is also no greater than ϵ_2 . Using this, we have that

$$\frac{1}{n} \sum_{i=1}^n |f_m(x_i) - \tilde{f}_m(x_i)| + \frac{1}{n} \sum_{i=1}^n |f_m(x'_i) - \tilde{f}_m(x'_i)| \leq 2v(1 + M_1/m_0)\epsilon_2.$$

The last conclusion about the cardinality of $\tilde{\mathcal{F}}$ follows from Lemma 10. \square

LEMMA 6. Let $\underline{x} = \{x_i\}_{i=1}^n$, where each x_i is a d -dimensional vector in \mathbb{R}^d . Define $\|\underline{x}\|_\infty^2 = \frac{1}{n} \sum_{i=1}^n \|x_i\|_\infty^2$. There is a subset $\tilde{\mathcal{H}}$ of \mathcal{H} with cardinality at most $\binom{2d+m}{m}$ such that for each $h(x) = \phi(x \cdot \theta)$ with $\|\theta\|_1 \leq \Lambda$ in \mathcal{H} , there is $\tilde{h}(x) = \phi(x \cdot \tilde{\theta})$ in $\tilde{\mathcal{H}}$ such that $\frac{1}{n} \sum_{i=1}^n |h(x_i) - \tilde{h}(x_i)|^2 \leq \Lambda \|\theta\|_1 \|\underline{x}\|_\infty^2 / m$.

PROOF. By the Lipschitz condition on ϕ , it is enough to prove the bound for $\frac{1}{n} \sum_{i=1}^n |\theta \cdot x_i - \tilde{\theta} \cdot x_i|^2$. Let v be a random vector that equals $e_j \text{sgn}(\theta_j) \Lambda$ with probability $|\theta_j|/\Lambda$, $j = 1, 2, \dots, d$ and equals the zero vector with probability $1 - \|\theta\|_1/\Lambda$. Let v_1, v_2, \dots, v_m be a random sample from the distribution defining v . Note that the average of $\bar{\theta} = \frac{1}{m} \sum_{j=1}^m v_j$ is θ and hence the average of each $|\bar{\theta} \cdot x_i - \theta \cdot x_i|^2$ is the variance of $v \cdot x_i$ divided by m . Taking the expectation of the desired quantity, we have

$$\begin{aligned} \mathbb{E} \frac{1}{n} \sum_{i=1}^n |\bar{\theta} \cdot x_i - \theta \cdot x_i|^2 &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} |\bar{\theta} \cdot x_i - \theta \cdot x_i|^2 \\ &\leq \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{E}(v \cdot x_i)^2}{m} \\ &\leq \frac{1}{n} \sum_{i=1}^n \frac{\|x_i\|_\infty^2 \Lambda \|\theta\|_1}{m} \\ &= \Lambda \|\theta\|_1 \|\underline{x}\|_\infty^2 / m. \end{aligned}$$

Since this bound holds on average, there must exist a realization $\tilde{\theta}$ of $\bar{\theta}$ for which the inequality is also satisfied. Consider the collection of all vectors of the form

$$(\Lambda/m) \sum_{j=1}^m u_j,$$

where u_j is any of the $2d + 1$ signed standard basis vectors including the zero vector. This collection has cardinality bounded by the number of non-negative integer solutions $q_1, q_2, \dots, q_{2d+1}$ to

$$q_1 + q_2 + \dots + q_{2d+1} = m.$$

This number is $\binom{2d+m}{m}$ with its logarithm is bounded by $m \log(e(2d/m + 1))$ or $2m \log(d + 1)$. An important aspect of the log cardinality of this empirical cover is that it is logarithmic (and not linear) in the dimension d . This

small dependence on d is what produces desirable risk bounds when d is significantly greater than the available sample size n . \square

LEMMA 7. *Let Z have mean zero and variance σ^2 . Moreover, suppose Z satisfies Bernstein's moment condition with parameter $\eta > 0$. Then*

$$(6.4) \quad \mathbb{E}(e^{tZ}) \leq \exp \left\{ \frac{t^2 \sigma^2 / 2}{1 - \eta |t|} \right\}, \quad |t| < 1/\eta.$$

LEMMA 8. *Define $Tf = \min\{B_n, |f|\} \operatorname{sgn} f$. Then*

- (I) $(y - Tf)^2 \leq (y - f)^2 + 2(|y| - B_n)^2 \mathbb{I}\{|y| > B_n\}$,
- (II) $(y - Tf)^2 \leq (y - T\tilde{f})^2 + 4B_n|f - \tilde{f}| + 4B_n(|y| - B_n) \mathbb{I}\{|y| > B_n\}$, and
- (III) $(T\tilde{f} - Tf)^2 \leq (f - f_1)^2 + 4B_n|f_1 - \tilde{f}|$.

PROOF. (I) Since $(y - Tf)^2 = (y - f)^2 + 2(f - Tf)(2y - f - Tf)$, the proof will be complete if we can show that

$$(f - Tf)(2y - f - Tf) \leq (|y| - B_n)^2 \mathbb{I}\{|y| > B_n\}.$$

Note that if $|f| \leq B_n$, the left hand side of the above expression is zero. Thus we may assume that $|f| > B_n$, in which case $f - Tf = \operatorname{sgn} f(|f| - B_n)$. Thus

$$\begin{aligned} (f - Tf)(2y - f - Tf) &= 2y \operatorname{sgn} f(|f| - B_n) - (|f| - B_n)(|f| + B_n) \\ &\leq 2|y|(|f| - B_n) - (|f| - B_n)(|f| + B_n). \end{aligned}$$

If $|y| \leq B_n$, the above expression is less than $-(|f| - B_n)^2 \leq 0$. Otherwise, it is a quadratic in $|f|$ that attains its global maximum at $|f| = |y|$. This yields a maximum value of $(|y| - B_n)^2$.

(II) For the second claim, note that

$$(y - Tf)^2 = (y - T\tilde{f})^2 + (T\tilde{f} - Tf)(2y - T\tilde{f} - Tf).$$

Hence, we are done if we can show that

$$(T\tilde{f} - Tf)(2y - T\tilde{f} - Tf) \leq 4B_n|f - \tilde{f}| + 4B_n(|y| - B_n) \mathbb{I}\{|y| > B_n\}.$$

If $|y| \leq B_n$, then

$$\begin{aligned} (T\tilde{f} - Tf)(2y - T\tilde{f} - Tf) &\leq 4B_n|T\tilde{f} - Tf| \\ &\leq 4B_n|\tilde{f} - f|. \end{aligned}$$

If $|y| > B_n$, then

$$\begin{aligned} (T\tilde{f} - Tf)(2y - T\tilde{f} - Tf) &\leq 2|T\tilde{f} - Tf||y| + 2B_n|T\tilde{f} - Tf| \\ &= 2|T\tilde{f} - Tf|(|y| - B_n) + 4B_n|T\tilde{f} - Tf| \\ &\leq 4B_n(|y| - B_n) + 4B_n|\tilde{f} - f|. \end{aligned}$$

(III) For the last claim, note that

$$\begin{aligned} (T\tilde{f} - Tf)^2 &= (T\tilde{f} - Tf_1)^2 + [2T\tilde{f} - Tf_1 - Tf](Tf_1 - Tf) \\ &\leq (T\tilde{f} - Tf_1)^2 + 4B_n|Tf_1 - Tf| \\ &\leq (\tilde{f} - f_1)^2 + 4B_n|f_1 - f| \end{aligned}$$

□

LEMMA 9. Let $Y = f^*(X) + \epsilon$ with $|f^*(X)| \leq B$. Suppose

(I) $\mathbb{E}e^{|\epsilon|/\nu} < +\infty$ or

(II) $\mathbb{E}e^{|\epsilon|^2/\nu} < +\infty$

for some $\nu > 0$. Then $\mathbb{E}[(Y^2 - B_n^2)\mathbb{I}\{|Y| > B_n\}]$ is at most

(I) $(4\nu^2/n)\mathbb{E}e^{|\epsilon|/\nu}$ provided $B_n > \sqrt{2}(B + \nu \log n)$ or

(II) $(2\nu/n)\mathbb{E}e^{|\epsilon|^2/\nu}$ provided $B_n > \sqrt{2}(B + \sqrt{\nu \log n})$.

PROOF. Under assumption (I),

$$\begin{aligned} \mathbb{P}(Y^2 - B_n^2 > t) &= \mathbb{P}(|Y| > \sqrt{t + B_n^2}) \\ &\leq \mathbb{P}(|\epsilon| > \sqrt{t + B_n^2} - B) \\ &\leq \mathbb{P}(|\epsilon| > (1/\sqrt{2})(\sqrt{t} + B_n) - B) \\ &\leq e^{-\frac{1}{\nu}\sqrt{\frac{t}{2}}} e^{-\frac{1}{\nu}(\frac{B_n}{\sqrt{2}} - B)} \mathbb{E}e^{|\epsilon|/\nu}. \end{aligned}$$

The last inequality follows from a simple application of Markov's inequality after exponentiation. Integrating the previous expression from $t = 0$ to $t = +\infty$ ($\int_0^\infty e^{-\frac{1}{\nu}\sqrt{\frac{t}{2}}} dt = 4\nu^2$) yields an upper bound on $\mathbb{E}[(Y^2 - B_n^2)\mathbb{I}\{|Y| > B_n\}]$ that is at most $(4\nu^2/n)\mathbb{E}e^{|\epsilon|/\nu}$ provided $B_n > \sqrt{2}(B + \nu \log n)$.

Under assumption (II),

$$\begin{aligned} \mathbb{P}(Y^2 - B_n^2 > t) &= \mathbb{P}(|Y|^2 > t + B_n^2) \\ &\leq \mathbb{P}(|\epsilon|^2 > (1/2)(t + B_n^2) - B^2) \\ &\leq e^{-\frac{t}{2\nu}} e^{-\frac{1}{\nu}(\frac{B_n^2}{2} - B^2)} \mathbb{E}e^{|\epsilon|^2/\nu}. \end{aligned}$$

The last inequality follows from a simple application of Markov's inequality after exponentiation. Integrating the previous expression from $t = 0$ to $t = +\infty$ ($\int_0^\infty e^{-\frac{t}{2\nu}} dt = 2\nu$) yields an upper bound on $\mathbb{E}[(Y^2 - B_n^2)\mathbb{I}\{|Y| > B_n\}]$ that is at most $(2\nu/n)\mathbb{E}e^{|\epsilon|^2/\nu}$ provided $B_n > \sqrt{2}(B + \sqrt{\nu \log n}) \geq \sqrt{2}(B^2 + \nu \log n)$. \square

LEMMA 10. *The number of functions having the form $\frac{\nu}{m} \sum_{k=1}^m h_k$, where h_k belong to a library of size M is at most $\binom{M-1+m}{m} \leq \binom{M+m}{m}$ with its logarithm bounded by $m \log(e(M/m + 1))$.*

PROOF. Suppose the elements in the library are indexed from 1 to M . Let q_i be the number of terms in $\sum_{k=1}^m h_k$ of type i . Hence the number of function of the form $\frac{\nu}{m} \sum_{k=1}^m h_k$ is at most the number of non-negative integer solutions q_1, q_2, \dots, q_M to $q_1 + q_2 + \dots + q_M = m$. This number is $\binom{M-1+m}{m}$ with its logarithm bounded by the minimum of $m \log(e((M-1)/m + 1))$ and $m \log M$. \square

Acknowledgements. The risk bound framework of this paper is self-contained, yet borrows heavily from the unpublished manuscript [15] written with Cong Huang and Gerald Cheang, which was judged to be too long for the Annals of Statistics. We are very indebted to their contributions. The present paper with Jason Klusowski manages a compression of much of that work, corrected determination of improved exponents of the rate, and new improved rates for the high dimensional setting. The authors would also like to thank William D. Brinda for useful discussions.

SUPPLEMENTARY MATERIAL

Supplement A: Proof of an approximation result and justification of (4.1).

(doi: [COMPLETED BY THE TYPESETTER](#); .pdf). We prove that a function satisfying a certain spectral condition belongs to the closure of the linear span of ridge ramp functions. We also give a justification for the penalty divided by sample size as stated in (4.1).

References.

- [1] ANTHONY, M. and BARTLETT, P. L. (1999). *Neural network learning: theoretical foundations*. Cambridge University Press, Cambridge. [MR1741038](#)

- [2] BARRON, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inform. Theory* **39** 930–945. [MR1237720](#)
- [3] BARRON, A. R. (1994). Approximation and estimation bounds for artificial neural networks. *Machine Learning* **14** 115–133.
- [4] BARRON, A., BIRGÉ, L. and MASSART, P. (1999). Risk bounds for model selection via penalization. *Probab. Theory Related Fields* **113** 301–413. [MR1679028](#)
- [5] BARRON, A. R. and KLUSOWSKI, J. M. (2016). Training artificial neural nets by adaptive annealing. *Preprint*.
- [6] BARRON, A. R., HUANG, C., LI, J. and LUO, X. (2008a). The MDL principle, penalized likelihoods, and statistical risk. *Workshop on Information Theory Methods in Science and Engineering, Tampere, Finland*.
- [7] BARRON, A. R., HUANG, C., LI, J. and LUO, X. (2008b). The MDL principle, penalized likelihoods, and statistical risk. *Festschrift in Honor of Jorma Rissanen on the Occasion of his 75th Birthday* Tampere University Press, Tampere, Finland. Editor Ioan Tabus.
- [8] BARRON, A. R., COHEN, A., DAHMEN, W. and DEVORE, R. A. (2008c). Approximation and learning by greedy algorithms. *Ann. Statist.* **36** 64–94. [MR2387964](#)
- [9] BARTLETT, P. L. (1998). The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Trans. Inform. Theory* **44** 525–536. [MR1607706](#)
- [10] BREIMAN, L. (1993). Hinging hyperplanes for regression, classification, and function approximation. *IEEE Trans. Inform. Theory* **39** 999–1013. [MR1237723](#)
- [11] BUNEA, F., TSYBAKOV, A. B. and WEGKAMP, M. H. (2007). Aggregation for Gaussian regression. *Ann. Statist.* **35** 1674–1697. [MR2351101](#)
- [12] CHEANG, G. L. H. (1998). Neural network approximation and estimation of functions. *Yale University, Department of Statistics PhD thesis*.
- [13] GAMMERMAN, A., VOVK, V. and VAPNIK, V. (1998). Learning by transduction. *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence* 148–155.
- [14] HUANG, C. (2008). Risk of penalized least squares, greedy selection and ℓ_1 -penalization for flexible function libraries. *Yale University, Department of Statistics PhD thesis*.
- [15] HUANG, C., CHEANG, G. L. and BARRON, A. R. (2008). Risk of penalized least squares, greedy selection and ℓ_1 penalization for flexible function libraries. *Yale University, Department of Statistics technical report*.
- [16] JONES, L. K. (1992). A simple lemma on greedy approximation in Hilbert space and convergence rates for projection pursuit regression and neural network training. *Ann. Statist.* **20** 608–613. [MR1150368](#)
- [17] JUDITSKY, A. and NEMIROVSKI, A. (2000). Functional aggregation for nonparametric regression. *Ann. Statist.* **28** 681–712. [MR1792783](#)
- [18] LEE, W. S., BARTLETT, P. L. and WILLIAMSON, R. C. (1996). Efficient agnostic learning of neural networks with bounded fan-in. *IEEE Trans. Inform. Theory* **42** 2118–2132. [MR1447518](#)
- [19] MAKOVOS, Y. (1996). Random approximants and neural networks. *J. Approx. Theory* **85** 98–109. [MR1382053](#)
- [20] YANG, Y. and BARRON, A. (1999). Information-theoretic determination of minimax rates of convergence. *Ann. Statist.* **27** 1564–1599. [MR1742500](#)
- [21] ZHANG, T. (2009). Some sharp performance bounds for least squares regression with L_1 regularization. *Ann. Statist.* **37** 2109–2144. [MR2543687](#)

**SUPPLEMENT TO "RISK BOUNDS FOR
HIGH-DIMENSIONAL RIDGE FUNCTION
COMBINATIONS INCLUDING NEURAL NETWORKS"**

BY JASON M. KLUSOWSKI AND ANDREW R. BARRON

Yale University

We prove that a function satisfying a certain spectral condition belongs to the closure of the linear span of ridge ramp functions. We also give a justification for the penalty divided by sample size as stated in (4.1).

THEOREM. *For an arbitrary target function $f^*(x) = \int_{\mathbb{R}^d} e^{ix \cdot \omega} \tilde{f}(\omega) d\omega$ with $v_{f^*,2} = \int_{\mathbb{R}^d} \|\omega\|_1^2 |\tilde{f}(\omega)| d\omega$ finite, there exists a linear combination of ridge ramp functions with $\|a_k\| = 1$ and $|b_k| \leq 1$ such that*

$$\|f^* - f_m\|^2 \leq \frac{16v_{f^*,2}^2}{m}.$$

PROOF. If f^* can be extended to a function on $L^2(\mathbb{R}^d)$ with Fourier transform \tilde{f} , the function $f^*(x) - x \cdot \nabla f^*(0) - f^*(0)$ can be written as the real part of

$$(6.5) \quad \int_{\mathbb{R}^d} (e^{i\omega \cdot x} - i\omega \cdot x - 1) \tilde{f}(\omega) d\omega.$$

If $|z| \leq c$, we note the identity

$$-\int_0^c [(z-u)_+ e^{iu} + (-z-u)_+ e^{-iu}] du = e^{iz} - iz - 1.$$

If $c = \|\omega\|_1$, $z = \omega \cdot x$, $\alpha = \alpha(\omega) = \omega/\|\omega\|_1$, and $u = \|\omega\|_1 t$, $0 \leq t \leq 1$, we find that

$$-\|\omega\|_1^2 \int_0^1 [(\alpha \cdot x - t)_+ e^{i\|\omega\|_1 t} + (-\alpha \cdot x - t)_+ e^{-i\|\omega\|_1 t}] dt = e^{i\omega \cdot x} - i\omega \cdot x - 1.$$

Primary 60K35, 60K35; secondary 60K35
Keywords and phrases: sample, L^AT_EX 2_ε

Multiplying the above by $\tilde{f}(\omega) = e^{ib(\omega)}|\tilde{f}(\omega)|$, integrating over \mathbb{R}^d , and applying Fubini's theorem yields

$$f^*(x) - x \cdot \nabla f^*(0) - f^*(0) = \int_{\mathbb{R}^d} \int_0^1 g(t, \omega) dt d\omega,$$

where

$$g(t, \omega) = -[(\alpha \cdot x - t)_+ \cos(\|\omega\|_1 t + b(\omega)) + (-\alpha \cdot x - t)_+ \cos(\|\omega\|_1 t - b(\omega))] \|\omega\|_1^2 |\tilde{f}(\omega)|.$$

Consider the density on $\{0, 1\} \times [0, 1] \times \mathbb{R}^d$ defined by

$$p(z, t, \omega) = (1 - \gamma)(1 - z)\mu(t, \omega) + \gamma z\nu(t, \omega)$$

where $\mu(t, \omega) = |\cos(\|\omega\|_1 t + b(\omega))| \|\omega\|_1^2 |\tilde{f}(\omega)| / U$ and $\nu(t, \omega) = |\cos(\|\omega\|_1 t - b(\omega))| \|\omega\|_1^2 |\tilde{f}(\omega)| / V$ with constants U and V chosen to make μ and ν integrate to one, respectively. Choose $\gamma = V / (U + V)$ and consider the random variable $\eta(z, t, \alpha)(x)$ that equals

$$(1 - z)(\alpha \cdot x - t)_+ s_\mu(t, \omega) + z(-\alpha \cdot x - t)_+ s_\nu(t, \omega),$$

where $s_\mu(t, \omega) = -\text{sgn} \cos(\|\omega\|_1 t + b(\omega))$ and $s_\nu(t, \omega) = -\text{sgn} \cos(\|\omega\|_1 t - b(\omega))$. Note that $\eta(z, t, \alpha)(x)$ has the form $\pm(\alpha \cdot x - t)_+$ and $B = U + V \leq 2v_{f^*, 2}$. Thus, we see that

$$f^*(x) - x \cdot \nabla f^*(0) - f^*(0) = B \int_{\{0, 1\} \times [0, 1] \times \mathbb{R}^d} \eta(z, t, \alpha)(x) dp(z \times t \times \omega).$$

One can obtain the final result by sampling $(z_1, t_1, \omega_1), \dots, (z_m, t_m, \omega_m)$ randomly from $p(z, t, \omega)$ and considering the average $\frac{1}{m} \sum_{k=1}^m h(z_k, t_k, \omega_k)$. Note that since $x = (x)_+ - (-x)_+$, we can regard $x \cdot \nabla f^*(0)$ as belonging to the linear span of $\{x \mapsto \pm(\alpha \cdot x - t)_+ : \|\alpha\|_1 = 1, |t| \leq 1\}$. An easy argument shows that its variance is bounded by $16v_{f^*, 2}^2/m$. This simple argument can be extended to higher order expansions of f^* . The function $f^*(x) - x^T H_{f^*}(0)x/2 - x \cdot \nabla f^*(0) - f^*(0)$ ($H_{f^*}(0)$ is the Hessian of f^* at the point zero) can be written as the real part of

$$(6.6) \quad \int_{\mathbb{R}^d} (e^{i\omega \cdot x} + (\omega \cdot x)^2/2 - i\omega \cdot x - 1) \tilde{f}(\omega) d\omega.$$

As before, the integrand in (6.6) admits an integral representation by

$$(i/2) \|\omega\|_1^3 \int_0^1 [(-\alpha \cdot x - t)_+^2 e^{-i\|\omega\|_1 t} - (\alpha \cdot x - t)_+^2 e^{i\|\omega\|_1 t}] dt.$$

Employing a sampling argument from an appropriately defined density, we are able to approximate $f^*(x) - x^T H_{f^*}(0)x/2 - x \cdot \nabla f^*(0)$ by a linear combinations of m second order spline functions (having bounded internal parameters) $(\alpha \cdot x - t)_+^2$ with a squared error bounded by $16v_{f^*}^2/3/m$. \square

In the next part of the supplement, we justify the form of the penalty divided by sample size as stated in (4.1). For this, we first need a lemma.

LEMMA. *Let $f = \sum_h \beta_h h$ and let $\tilde{\mathcal{H}}$ be an empirical L^2 ϵ -net of \mathcal{H} of size M . There is a choice of $\tilde{h}_1, \dots, \tilde{h}_m$ in $\tilde{\mathcal{H}}$ with $\tilde{f}_m = (v/m) \sum_{k=1}^m \tilde{h}_k$, $v \geq v_f$ such that*

$$P_n(T\tilde{f}||f^*) - 2P_n(f||f^*) + D'_n(\tilde{f}, f) \leq \frac{4vv_f}{m} + 6v^2\epsilon^2 + 4\sigma v\epsilon + T_n + R_n,$$

where $R_n = \frac{1}{n} \sum_{i=1}^n (\epsilon_i^2 - \sigma^2) \mathbb{I}\{\frac{1}{n} \sum_{i=1}^n (\epsilon_i^2 - \sigma^2) > 0\}$. Moreover, $\mathbb{E}R_n \leq \frac{4\sqrt{3}\eta\sigma}{\sqrt{n}}$ and if $\tilde{\mathcal{F}}$ denotes the collection of functions of the form \tilde{f}_m , then $\tilde{\mathcal{F}}$ has cardinality at most $\binom{M+m}{m}$.

PROOF. Let $f_m = (v/m) \sum_{k=1}^m h_k$ be as in Lemma 3 and $\tilde{f} = \tilde{f}_m = (v/m) \sum_{k=1}^m \tilde{h}_k$ be as in Lemma 5. Furthermore, define

$$\tau_n = \min\{1, \sqrt{D_n(\tilde{f}_m, f_m)/\sigma^2}\}.$$

By Lemma 8 (I),

$$P_n(T\tilde{f}_m||f^*) \leq P_n(\tilde{f}_m||f^*) + T_n,$$

and hence

$$P_n(T\tilde{f}_m||f^*) - (1 + \tau_n)P_n(f||f^*) \leq P_n(\tilde{f}_m||f^*) - (1 + \tau_n)P_n(f||f^*) + T_n.$$

Also observe that

$$P_n(\tilde{f}_m||f^*) - (1 + \tau_n)P_n(f||f^*) = [P_n(\tilde{f}_m||f^*) - (1 + \tau_n)P_n(f_m||f^*)] + (1 + \tau_n)[P_n(f_m||f^*) - P_n(f||f^*)],$$

and

$$D'_n(\tilde{f}_m, f) \leq 2D'_n(\tilde{f}_m, f_m) + 2D'_n(f, f_m).$$

By Lemma 3,

$$P_n(f_m||f^*) - P_n(f||f^*) + D'_n(\tilde{f}, f) \leq \frac{2vv_f}{m}.$$

Using the inequality $a^2 - (1 + \tau)b^2 \leq \frac{1+\tau}{\tau}(a-b)^2$, we have

$$\begin{aligned} P_n(\tilde{f}_m || f^*) - (1 + \tau_n)P_n(f_m || f^*) &\leq \frac{1 + \tau_n}{\tau_n} D_n(\tilde{f}_m, f_m) + \tau_n \bar{\epsilon}_n^2 \\ &\leq \frac{1 + \tau_n}{\tau_n} D_n(\tilde{f}_m, f_m) + \tau_n \sigma^2 + R_n \\ &\leq 3D_n(\tilde{f}_m, f_m) + 2\sigma \sqrt{D_n(\tilde{f}_m, f_m)} + R_n, \end{aligned}$$

where the last inequality follows from the definition of τ_n and $\bar{\epsilon}_n^2 = \frac{1}{n} \sum_{i=1}^n \epsilon_i^2$. The first conclusion is proved after observing that $D_n(\tilde{f}_m, f_m) \leq 2v^2 \epsilon^2$.

The statement about the cardinality of $\tilde{\mathcal{F}}$ follows immediately from the form of \tilde{f} and Lemma 10.

We are also able to comment on how R_n behaves on average. To see this, note that $R_n \leq s + R_n \mathbb{I}\{R_n > s\}$ for all $s > 0$. Now, $R_n \mathbb{I}\{R_n > s\} = \frac{1}{n} \sum_{i=1}^n (\epsilon_i^2 - \sigma^2) \mathbb{I}\{\frac{1}{n} \sum_{i=1}^n (\epsilon_i^2 - \sigma^2) > s\}$. By Chebychev's inequality and the fact that $\mathbb{V}(\epsilon) \leq \sigma^2$ and $\mathbb{V}(\epsilon^2) \leq 12\eta^2 \sigma^2$ (as per the Bernstein moment condition on the noise), $\mathbb{P}(\frac{1}{n} \sum_{i=1}^n (\epsilon_i^2 - \sigma^2) > s + t) \leq \frac{\mathbb{V}(\epsilon^2)}{n(t+s)^2}$. Integrating this inequality from $t = 0$ to $t = +\infty$ yields $R_n \mathbb{I}\{R_n > s\} \leq \frac{\mathbb{V}(\epsilon^2)}{ns}$. Choosing $s = \sqrt{\mathbb{V}(\epsilon^2)/n}$ produces $\mathbb{E}R_n \leq \frac{2\sqrt{\mathbb{V}(\epsilon^2)}}{\sqrt{n}} \leq \frac{4\sqrt{3}\eta\sigma}{\sqrt{n}}$. \square

By the above lemma and similar arguments leading to the conclusion of Theorem 3, we see that a valid penalty divided by sample size is at least

$$\frac{\gamma_n m \log M}{n} + \frac{4vv_f}{m} + 6v^2 \epsilon^2 + 2\sigma v \epsilon + T_n + R_n$$

By Lemma 6, an empirical L^2 ϵ -cover of \mathcal{H} has cardinality less than $\binom{2d + \lceil (\Lambda/\epsilon)^2 \rceil}{\lceil (\Lambda/\epsilon)^2 \rceil}$.

The logarithm of $\binom{2d + \lceil (\Lambda/\epsilon)^2 \rceil}{\lceil (\Lambda/\epsilon)^2 \rceil}$ is bounded by $4(\Lambda/\epsilon)^2 \log(d+1)$. Thus we see that a valid penalty divided by sample size is at least

$$(6.7) \quad \frac{4\gamma_n m (\Lambda/\epsilon)^2 \log(d+1)}{n} + \frac{4vv_f}{m} + 6v^2 \epsilon^2 + 4\sigma v \epsilon + T_n + R_n$$

Choose m to be the ceiling of $\left(\frac{nvv_f \epsilon^2}{\gamma_n \Gamma^2 \log(d+1)}\right)^{1/2}$. Plugging this back into (6.7), we find that a valid penalty divided by sample size is at least

$$8 \left(\frac{vv_f \gamma_n \Lambda^2 \log(d+1)}{n \epsilon^2} \right)^{1/2} + \frac{4\gamma_n \Lambda^2 \log(d+1)}{n \epsilon^2} + 6v^2 \epsilon^2 + 4\sigma v \epsilon + T_n + R_n.$$

Choose $v = v_f$ and

$$\epsilon = \frac{1}{\left(\frac{v^2 n}{\gamma_n \Lambda^2 \log(d+1)}\right)^{1/6} + \left(\frac{\sigma^2 n}{\gamma_n \Lambda^2 \log(d+1)}\right)^{1/4}}.$$

With this choice of ϵ , a valid penalty divided by sample size is seen to be at least a modest constant multiple of

$$\begin{aligned} & \left(\frac{v_f^4 \Lambda^2 \gamma_n \log(d+1)}{n}\right)^{1/3} + \sqrt{\sigma} \left(\frac{v_f^4 \Lambda^2 \gamma_n \log(d+1)}{n}\right)^{1/4} + \\ & \left(\frac{v_f^2 \Lambda^2 \gamma_n \log(d+1)}{n}\right)^{2/3} + \left(\frac{\Lambda^2 \gamma_n \log(d+1)}{n}\right)^{2/3} + \\ & \sqrt{\sigma} \left(\frac{\Lambda^{4/3} \gamma_n \log(d+1)}{n}\right)^{3/4} + T_n + R_n. \end{aligned}$$

Thus when $\frac{\Lambda^2 \gamma_n \log(d+1)}{n}$ is small, we see that the penalty divided by sample size has the form (4.1) plus negligible terms that do not depend on the candidate fit.

YALE UNIVERSITY
DEPARTMENT OF STATISTICS
24 HILLHOUSE AVENUE
NEW HAVEN, CONNECTICUT
E-MAIL: andrew.barron@yale.edu
jason.klusowski@yale.edu