

On the Multiple Descent of Minimum-Norm Interpolants and Restricted Lower Isometry of Kernels

Tengyuan Liang

University of Chicago, Booth School of Business

Alexander Rakhlin

MIT

Xiyu Zhai

MIT

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We study the risk of minimum-norm interpolants of data in Reproducing Kernel Hilbert Spaces. Our upper bounds on the risk are of a *multiple-descent* shape for the various scalings of $d = n^\alpha$, $\alpha \in (0, 1)$, for the input dimension d and sample size n . Empirical evidence supports our finding that minimum-norm interpolants in RKHS can exhibit this unusual non-monotonicity in sample size; furthermore, locations of the peaks in our experiments match our theoretical predictions. Since gradient flow on appropriately initialized wide neural networks converges to a minimum-norm interpolant with respect to a certain kernel, our analysis also yields novel estimation and generalization guarantees for these over-parametrized models.

At the heart of our analysis is a study of spectral properties of the random kernel matrix restricted to a filtration of eigen-spaces of the population covariance operator, and may be of independent interest.

Keywords: interpolation; kernels; consistency; generalization

1. Introduction

We investigate the generalization and consistency of minimum-norm interpolants

$$\hat{f} \in \operatorname{argmin}_{f \in \mathcal{H}} \|f\|_{\mathcal{H}} \quad \text{s.t.} \quad f(x_i) = y_i, \quad i = 1, \dots, n \quad (1)$$

of the data $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^d \times \mathbb{R}$ with respect to a norm in a Reproducing Kernel Hilbert Space \mathcal{H} . The interpolant, also termed “Kernel Ridgeless Regression,” can be viewed as a limiting solution of

$$\operatorname{argmin}_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2 \quad (2)$$

as $\lambda \rightarrow 0$. Classical statistical analyses of Kernel Ridge Regression (see e.g. [Caponnetto and De Vito \(2007\)](#) and references therein) rely on a carefully chosen regularization parameter to control the bias-variance tradeoff, and the question of consistency of the non-regularized solution falls outside the scope of these classical results.

Recent literature has focused on understanding risk of estimators that interpolate data, including work on nonparametric local rules (Belkin et al., 2018b,d), high-dimensional linear regression (Bartlett et al., 2019; Hastie et al., 2019), random features model (Ghorbani et al., 2019), classification with rare instances (Feldman, 2019), and kernel (ridgeless) regression (Belkin et al., 2018c; Liang and Rakhlin, 2018; Rakhlin and Zhai, 2018).

This paper continues the line of work on kernel regression. More precisely, Rakhlin and Zhai (2018) showed that the minimum-norm interpolant with respect to Laplace kernel is *not* consistent (that is, risk does not go to zero with $n \rightarrow \infty$) if dimensionality d of the data is constant with respect to n , even if the bandwidth of the kernel is chosen adaptively. On the other hand, Liang and Rakhlin (2018) investigated the regime $n \asymp d$ and showed that risk can be upper bounded by a quantity that can be small under favorable spectral properties of the data and the kernel matrix.

The present paper aims to paint a more comprehensive picture, studying the performance of the minimum-norm interpolants in a general scaling regime $d \asymp n^\alpha$, $\alpha \in (0, 1)$. Figure 1 summarizes the non-monotone behavior of our upper bound on the risk of the minimum-norm interpolant, as reported in Theorem 1 below.




Figure 1: **Multiple-descent behavior** of the rates as the scaling $d = n^\alpha$ changes.

We make two observations. First, for any integer $\iota \geq 1$, for $\alpha \in [\frac{1}{\iota+1}, \frac{1}{\iota})$, there exists a “valley” on the curve at each $d = n^{\frac{1}{\iota+1/2}}$ where the rate is fast (of the order $n^{-\beta}$ with $\beta = \frac{1}{2\iota+1}$). Second, towards the lower-dimensional regime (α moving towards 0), the fastest possible rate even at the bottom of the valley is getting worse, with no consistency in the $\alpha = 0$ regime, matching the lower bound of (Rakhlin and Zhai, 2018).

Depending on the point of view, we can also interpret the upper bound of Theorem 1 by fixing d and analyzing the behavior in n . In this case, the interpretation is rather counterintuitive: more data can lead to alternating regimes of better and worse performance. Conceptually, this occurs because with more samples, the empirical kernel matrix could estimate more complex subspaces associated with smaller population eigenvalues, thus increasing the variance of the interpolants.

Our experiments, reported in Figure 2, confirm the surprising multiple-peaks behavior, suggesting that the non-monotone shape of our upper bound is not just an artifact of the proof technique. Moreover, the locations of the peaks line up with our theoretical predictions. Our finding complements the *double-descent* behavior investigated previously in the literature (Belkin et al., 2018a;

Mei and Montanari, 2019),¹ suggesting that the behavior in the kernel case is significantly more detailed.

The challenging problem of proving a *lower bound* that exhibits the multiple descent behavior remains open. A more detailed analysis that studies relative heights of the peaks also appears to be an interesting direction of investigation. While the peaks and their size are certainly interesting, the reader should also note the positive message of our main result: the interpolating solution provably has a diminishing (in sample size) out-of-sample error for most of the scalings of d and n .




Figure 2: Empirical evidence for multiple descent. Sample size $n = 5000$, x-axis: $\log d$, y-axis: variance of the minimum-norm interpolant. Vertical lines denote theoretically predicted peaks, calculated by solving for d in the equation $\binom{d+i}{i} = n$ with $i = 1, 2, 3, 4$. Here we use the exponential kernel $h(t) = \exp(t)$, with $x_i \sim \mathcal{N}(0, I_d)$ and $y_i \sim \mathcal{N}(0, 1)$ sampled independently. In such a case, the risk only contains the variance term.

The main result of the paper can be informally stated as follows.

Theorem 1 (Informal) *For any integer $\iota \geq 1$, consider $d = n^\alpha$ where $\alpha \in [\frac{1}{\iota+1}, \frac{1}{\iota})$. Consider a general function $h \in C^\infty(\mathbb{R})$ and define the inner product kernel $k(x, z) = h(x^\top z/d)$. Consider n i.i.d. data pairs (x_i, y_i) drawn from $\mathcal{P}_{X,Y}$, and denote the target function $f_*(x) = \mathbb{E}[Y|X = x]$. Suppose the conditions on f_* , h and $\mathcal{P}_{X,Y}$ specified by Theorems 2-3 are satisfied. With probability at least $1 - \delta - e^{-n/d^\iota}$ on the design $\mathbf{X} \in \mathbb{R}^{n \times d}$,*

$$\mathbb{E} \left[\|\hat{f} - f_*\|_{\mathcal{P}_X}^2 | \mathbf{X} \right] \leq C \cdot \left(\frac{d^\iota}{n} + \frac{n}{d^{\iota+1}} \right) \asymp n^{-\beta},$$

$$\beta := \min \{ (\iota + 1)\alpha - 1, 1 - \iota\alpha \} \geq 0.$$

Here the constant $C(\delta, h, \iota, \mathcal{P})$ does not depend on d, n , with \mathcal{P} denoting the distribution of each coordinate of X .

It is easy to see that the minimum-norm interpolant in (1) has the closed-form solution

$$\hat{f}(x) = k(x, \mathbf{X})^\top K^{-1} \mathbf{Y}$$

1. In Figure 2, we only plotted the variance of the minimum-norm interpolant since the shape will dominate the bias term for appropriately scaled conditional variance of the Y variable.

if the kernel matrix $K \in \mathbb{R}^{n \times n}$ is invertible. Here $k(x, \mathbf{X}) = [k(x, x_1), \dots, k(x, x_n)]^\top$, $K_{i,j} = k(x_i, x_j)$, and $\mathbf{Y} = [y_1, \dots, y_n]^\top$. As discussed below, the variance of the estimator can be upper bounded by

$$\sigma_Y^2 \cdot \mathbb{E}_{x, \mathbf{X}} \|k(x, \mathbf{X})^\top K^{-1}\|^2$$

where σ_Y^2 is a uniform upper bound on the conditional variance of Y given X . However, in general, the smallest eigenvalue of the sample kernel matrix scales as a constant. Hence, further estimates on the variance term require a careful spectral analysis of the sample-based K and the population-based $k(x, \mathbf{X})$. Note that the eigenvalues of the empirical kernel matrix have one-to-one correspondance to that of the empirical covariance operator. We prove that on a filtration of eigen-spaces of the covariance operator defined by the population distribution, the empirical covariance operator satisfies a certain *restricted lower isometry* property. This spectral analysis is the main technical part.

2. Restricted Lower Isometry of High-Dimensional Kernels

In this section, we highlight Proposition 1, which establishes the *Restricted Lower Isometry Property*. This property proves crucial in bounding the generalization error for the kernel ridgeless regression and, as a consequence, for randomly-initialized wide neural networks trained to convergence. The detailed proof of Proposition 1 is deferred to Section A.

2.1. Setup

Before stating the main proposition, let us introduce the formal setup and assumptions for the rest of the paper. Random vectors $x_1, \dots, x_n \in \mathbb{R}^d$ are drawn i.i.d. from a product distribution $\mathcal{P}_X = \mathcal{P}^{\otimes d}$, where the distribution for each coordinate \mathcal{P} is independent and satisfies the following property.

Assumption 1 (Distribution for each coordinate) *Assume that: (1) $\mathbb{E}_{z \sim \mathcal{P}}[z] = 0$ and for a constant $\nu > 1$, $\mathbb{P}(|z| \geq t) \leq C(1+t)^{-\nu}$ holds for all $t \geq 0$. (2) For any set S of finitely many real numbers, $\mathbb{P}(z \in S) < 1$.*

In addition, we require that $\forall x \in \mathcal{X} \subseteq \mathbb{R}^d$, the conditional variance is bounded by a constant: $\text{Var}[Y|X=x] \leq \sigma_Y^2$.

Consider a function $h \in \mathcal{C}^\infty(\mathbb{R})$ whose Taylor expansion converges for all $t \in \mathbb{R}$

$$h(t) = \sum_{i=0}^{\infty} \alpha_i t^i \tag{3}$$

with all coefficients $\alpha_i \geq 0$. We define a kernel function $k(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ induced by h

$$k(x, z) := h\left(\frac{x^\top z}{d}\right). \tag{4}$$

Similarly, we define the *normalized* finite dimensional kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$,

$$\mathbf{K}_{ij} := k(x_i, x_j)/n, \quad 1 \leq i, j \leq n$$

In other words, $\mathbf{K} = k(\mathbf{X}, \mathbf{X})/n$ with the $1/n$ normalization.

Denote the truncated polynomial and the corresponding truncated kernel matrix $\mathbf{K}^{[\leq \iota]}$ (used only in the proof) as

$$h^{[\leq \iota]}(t) = \sum_{i=0}^{\iota} \alpha_i t^i, \quad \mathbf{K}_{ij}^{[\leq \iota]} := h^{[\leq \iota]}(x_i^\top x_j/d)/n.$$

Similarly, we denote the degree- ι component and the corresponding kernel as

$$h^{[\iota]}(t) = \alpha_\iota t^\iota, \quad \mathbf{K}_{ij}^{[\iota]} := h^{[\iota]}(x_i^\top x_j/d)/n.$$

We are interested in the following high dimensional regime: there is a fixed positive integer ι such that

$$\iota < \frac{\log n}{\log d(n)} < \iota + 1. \quad (5)$$

Our investigation focuses on the regime when the dimension $d(n)$ grows with the sample size n , with n being sufficiently large.

Finally, we use the norm $\|g\|_{L_2(\mathcal{P}_X)}^2 = \|g\|_{\mathcal{P}_X}^2 = \int g^2 \mathcal{P}_X(dx)$ to measure the quality of the estimator. Thanks to the identity $\mathbb{E} \left\| \hat{f} - f_* \right\|_{\mathcal{P}_X}^2 = \mathbb{E}(\hat{f}(X) - Y)^2 - \min_{f \in \mathcal{H}} \mathbb{E}(f(X) - Y)^2$, our results on estimation directly translate into prediction error guarantees.

2.2. Main Technical Result

Now we are ready to state the main technical contribution. We establish the restricted lower isometry property of the empirical kernel matrix on a filtration of spaces indexed by the polynomial basis with increasing degree.

Proposition 1 (Restricted Lower Isometry of High Dimensional Kernel) *Let ι_0 be a positive integer. Assume that the first $\iota_0 + 1$ Taylor coefficients $\alpha_0, \dots, \alpha_{\iota_0}$ of the function h (defined in (3)) are positive. Assume that $d^{\iota_0} \log d = o(n)$. Let Assumption 1 on \mathcal{P} be satisfied with $\nu > \iota_0$.*

Then there are positive constants C, C' depending only on ι_0, \mathcal{P} , and $\{\alpha_i\}_{i \leq \iota_0}$ such that for n large enough, with probability at least $1 - e^{-\Omega(n/d^{\iota_0})}$ the following holds:

- for any $\iota \leq \iota_0$, $\mathbf{K}^{[\leq \iota]}$ has $\binom{\iota+d}{\iota}$ nonzero eigenvalues, all of them larger than $C' d^{-\iota}$, and
- the range of $\mathbf{K}^{[\leq \iota]}$ is

$$\text{span} \left\{ (p(x_1), \dots, p(x_n)) : p \text{ is a multivariable polynomial of degree not larger than } \iota \right\}.$$

2.3. Proof Outline.

First, observe that

$$n\mathbf{K}_{ij} = \sum_{\iota=0}^{\infty} \alpha_\iota \left(\frac{x_i^\top x_j}{d} \right)^\iota = \sum_{r_1, \dots, r_d \geq 0} c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d} p_{r_1 \dots r_d}(x_i) p_{r_1 \dots r_d}(x_j) / d^{r_1 + \dots + r_d}$$

with $c_{r_1 \dots r_d} = \frac{(r_1 + \dots + r_d)!}{r_1! \dots r_d!}$ and monomials $p_{r_1 \dots r_d}(x_i) = (x_i[1])^{r_1} \dots (x_i[d])^{r_d}$ with multi-index $r_1 \dots r_d$. The degree-bounded empirical kernel is then

$$n\mathbf{K}_{ij}^{[\leq \iota]} := \sum_{\substack{r_1, \dots, r_d \geq 0 \\ r_1 + \dots + r_d \leq \iota}} c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d} p_{r_1 \dots r_d}(x_i) p_{r_1 \dots r_d}(x_j) / d^{r_1 + \dots + r_d} = \Phi \cdot \Phi^\top$$

with polynomial features $\Phi \in \mathbb{R}^{n \times \binom{\iota+d}{\iota}}$ of the form

$$\Phi_{i,(r_1 \dots r_d)} = (c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d})^{1/2} p_{r_1 \dots r_d}(x_i) / d^{(r_1 + \dots + r_d)/2} .$$

The restricted lower isometry of the kernel is equivalent to establishing that all eigenvalues of the sample covariance operator

$$\Theta^{[\leq \iota]} := \frac{1}{n} \Phi^\top \cdot \Phi$$

are lower bounded by $d^{-\iota}$. Observe that non-zero eigenvalues of $\mathbf{K}^{[\leq \iota]}$ have one-to-one correspondence to that of $\Theta^{[\leq \iota]}$.

Hypothetically, if the monomials $p_{r_1 \dots r_d}(x)$ were orthogonal in $L^2_{\mathcal{P}_X}$, then we would have

$$\mathbb{E} \left[\Theta^{[\leq \iota]} \right] = \text{diag}(C(0), \dots, C(\iota') \cdot d^{-\iota'}, \dots, \overbrace{C(\iota) \cdot d^{-\iota}}^{(\binom{d+\iota-1}{d-1} \text{ such entries})})$$

which would prove what we aim to establish on the smallest eigenvalue, at least in expectation. However, the orthogonality does not hold, and the monomials have a complex covariance structure that we have to tackle. To address the problem, we perform the *Gram-Schmidt* process on polynomials

$$\{1, t, t^2, \dots\} \rightarrow \{1, q_1(t), q_2(t), \dots\} \quad q \text{ orthogonal polynomial basis on } L^2_{\mathcal{P}},$$

and show that such basis is *weakly-correlated*. Then under the new polynomial features

$$\Phi_{i,(r_1 \dots r_d)} \rightarrow \Psi_{i,(r_1 \dots r_d)} = (c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d})^{1/2} \prod_{j \in [d]} q_{r_j}(x_i[j]) / d^{(r_1 + \dots + r_d)/2}$$

$$\Phi = \Psi \Lambda, \quad \Lambda \in \mathbb{R}^{\binom{\iota+d}{\iota} \times \binom{\iota+d}{\iota}} \text{ upper-triangular} .$$

It turns out that through technical calculations, we can show that such *weakly-correlated* polynomial features ensure that

$$\|\Lambda\|_{\text{op}}, \|\Lambda^{-1}\|_{\text{op}} \leq C(\iota) .$$

We can now focus on studying the smallest eigenvalue of the un-correlated features since for any $u \in \mathbb{R}^{\binom{\iota+d}{\iota}}$,

$$u^\top \Theta^{[\leq \iota]} u = \frac{1}{n} \|\Phi u\|^2 = \frac{1}{n} \|\Psi \Lambda u\|^2 \geq \lambda_{\min} \left(\frac{1}{n} \Psi^\top \Psi \right) \|\Lambda u\|^2 \asymp \lambda_{\min} \left(\frac{1}{n} \Psi^\top \Psi \right) \|u\|^2 .$$

The next challenge is in establishing a lower bound on the above smallest eigenvalue. Here, a naive use of standard concentration fails to provide strong high probability bounds. To see this, recall that if one wants to establish deviation bound via standard concentration like below

$$\sup_{u \in B_2^{\binom{d+\iota}{\iota}}} u^\top \left(\frac{1}{n} \Psi^\top \Psi - \mathbb{E} \left[\frac{1}{n} \Psi^\top \Psi \right] \right) u \lesssim \frac{\text{complexity}(B_2^{\binom{d+\iota}{\iota}})}{\sqrt{n}} ,$$

the deviation bound will typically be larger than $d^{-\iota}$ for ι of our interest. To address this, we take the *small-ball* approach, pioneered in [Koltchinskii and Mendelson \(2015\)](#); [Mendelson \(2014\)](#), which utilizes the non-negativity of the quadratic process. The intuition is as follows: due to positivity of $\langle \Psi(x_i), u \rangle^2$, the following lower bound holds

$$\|\Psi u\|^2 = \frac{1}{n} \sum_{i=1}^n \langle \Psi(x_i), u \rangle^2 \geq c_1 \mathbb{E}[\langle \Psi(x_i), u \rangle^2] \cdot \frac{1}{n} \sum_{i=1}^n I_{\langle \Psi(x_i), u \rangle^2 \geq c_1 \mathbb{E}[\langle \Psi(X), u \rangle^2]} .$$

Suppose one can show that there exist absolute constants $c_1, c_2 > 0$ such that

$$\mathbb{P}(\langle \Psi(x_i), u \rangle^2 \geq c_1 \mathbb{E}[\langle \Psi(X), u \rangle^2]) \geq c_2 ,$$

a condition referred to as the *small-ball* property. Then it immediately follows that with probability at least $1 - \exp(-c \cdot n)$

$$\frac{1}{n} \sum_{i=1}^n I_{\langle \Psi(x_i), u \rangle^2 \geq c_1 \mathbb{E}[\langle \Psi(X), u \rangle^2]} \geq c_2/2 .$$

Now the union bound on $B_2^{\binom{d+\iota}{\iota}}$ does not affect the rate significantly since the probability control is overwhelmingly small (exponential in n). Last but not least, the technicality remains to verify the *small-ball* property for weakly dependent polynomials via *Paley-Zygmund* inequality.

We note that concurrent work of [Ghorbani et al. \(2019\)](#) also implies a similar control on the least eigenvalue under a different setting with different assumptions on the underlying data. Specifically, their result concerns the approximation error on random Fourier feature models. It could be translated to a risk bound due to the dual relationship between random features and random samples, in the case when there is no label noise $y_i = f_*(x_i)$.

The rest of the paper is organized as follows. In Section 3 and 4 we will apply the key Proposition 1 to obtain generalization results for kernel ridgeless regression and wide neural networks, respectively. Sections 5 and 6 will be devoted to the proofs of Theorem 2 and Theorem 3 on the variance and bias of the minimum-norm interpolant. Section A in the Appendix will focus on the main steps behind proving Proposition 1. Appendix also contains several supporting lemmas.

3. Application to Kernel Ridgeless Regression

The following bias-variance decomposition holds, conditionally on \mathbf{X} :

$$\mathbb{E}_{\mathbf{Y}} \left\| \hat{f} - f_* \right\|_{\mathcal{P}_X}^2 = \mathbb{E}_{\mathbf{Y}} \left\| \hat{f} - \mathbb{E}_{\mathbf{Y}}[\hat{f}] \right\|_{\mathcal{P}_X}^2 + \left\| \mathbb{E}_{\mathbf{Y}}[\hat{f}] - f_* \right\|_{\mathcal{P}_X}^2 . \quad (6)$$

Here $\mathbb{E}_{\mathbf{Y}}$ denotes expectation only over the \mathbf{Y} vector. In this section, we refer to the first term as *Variance*, and the second term as (squared) *Bias*. As mentioned earlier, the variance term can be upper bounded as

$$\sigma_Y^2 \cdot \mathbb{E}_{x \sim \mathcal{P}_X} \left\| k(x, \mathbf{X})^\top k(\mathbf{X}, \mathbf{X})^{-1} \right\|^2 ,$$

thanks to the closed-form of \hat{f} . In the rest of this section we shall assume, for brevity, that $\sigma_Y^2 = 1$.

3.1. Variance

In this section, to control the variance term we make a stronger assumption on the tail behavior of \mathcal{P} .

Theorem 2 (Variance) *Let $x \sim \mathcal{P}_X$, $\mathbf{X} \sim \mathcal{P}^{n \times d}$, and \mathcal{P} be sub-Gaussian. Consider $h \in \mathcal{C}^\infty(\mathbb{R})$, denote $h(t) = \sum_{i=0}^{\infty} \alpha_i t^i$ with corresponding Taylor coefficients $\{\alpha_i\}_{i=0}^{\infty}$. Consider $k(x, z) = h(x^T z/d)$ for $x, z \in \mathcal{X}$.*

(i) *Suppose that:*

- $\alpha_1, \dots, \alpha_\iota > 0$;
- *there is $\iota' \geq 2\iota + 3$ such that $\alpha_{\iota'} > 0$.*

Assume $d^\iota \log d \lesssim n \lesssim d^{\iota+1}$. Then with probability at least $1 - e^{-\Omega(n/d^\iota)}$ w.r.t. \mathbf{X} ,

$$\text{Variance} \leq \mathbb{E}_{x \sim \mathcal{P}_X} \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 \leq C \left(\frac{d^\iota}{n} + \frac{n}{d^{\iota+1}} \right). \quad (7)$$

(ii) *Suppose that the Taylor expansion coefficients satisfy for some $\iota > 0$:*

- $\alpha_1, \dots, \alpha_\iota > 0$;
- $\forall \iota' > \iota, \alpha_{\iota'} = 0$, i.e. *k is a polynomial kernel.*

Assume $d^\iota \log d \lesssim n$. Then with probability at least $1 - e^{-\Omega(n/d^\iota)}$ w.r.t. \mathbf{X} ,

$$\text{Variance} \leq \mathbb{E}_{x \sim \mathcal{P}_X} \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 \leq C \frac{d^\iota}{n}. \quad (8)$$

The proof of the above theorem, which appears in Section 5, depends on breaking the variance into two parts depending on the polynomial degree: the first part we can further upper bounded via the key restricted lower isometry proposition on a filtration of spaces (ordered according to the degree of the polynomials), and for the second part we utilize the fact that $n\mathbf{K} \succeq \Omega(1)\mathbf{I}_n$.

3.2. Bias

In this section, we bound the bias part for the min-norm interpolated solution. In fact, we will show that, under a suitable assumption, the squared bias is upper bounded by a multiple factor of the variance term, studied in the previous section.

Theorem 3 (Bias) *Assume that the target function $f_*(x) = \mathbb{E}[Y|X = x]$ can be represented as*

$$f_*(x) = \int_{\mathcal{X}} k(x, z) \rho_*(z) \mathcal{P}_X(dz) \quad (9)$$

with $\int_{\mathcal{X}} \rho_^4(x) \mathcal{P}_X(dx) \leq C$ for $C > 0$. Assume that $\sup_{x \in \mathcal{X}} k(x, x) \leq M$. Then we have*

$$\begin{aligned} \text{Bias} &:= \mathbb{E}_{x \sim \mathcal{P}_X} \left(k(x, \mathbf{X}) k(\mathbf{X}, \mathbf{X})^{-1} f_*(\mathbf{X}) - f_*(x) \right)^2 \\ &\leq C_1(\mathbf{X}) \cdot \mathbb{E}_x \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 + \frac{C_2(\mathbf{X})}{n} \end{aligned}$$

where the scalar random variables are bounded in ℓ_2 -sense: $\mathbb{E}_{\mathbf{X}}[C_1(\mathbf{X})]^2, \mathbb{E}_{\mathbf{X}}[C_2(\mathbf{X})]^2 \lesssim 1$.

We remark that $\mathbb{E}_x \|k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, x)\|^2$ is the expression for the upper bound on variance in Section 3.1. The statement can be strengthened to the ‘‘in probability’’ statement, as follows

$$\begin{aligned} \text{Bias} &:= \|k(x, \mathbf{X})k(\mathbf{X}, \mathbf{X})^{-1}f_*(\mathbf{X}) - f_*(x)\|_{\mathcal{P}_X}^2 \\ &\leq \frac{1}{\sqrt{\delta}} \cdot \left(\mathbb{E}_x \|k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, x)\|^2 \vee \frac{1}{n} \right) \end{aligned}$$

with probability $1 - \delta$ on \mathbf{X} . See Proposition 4 for details. We emphasize that here the factor $\delta^{-1/2}$ has *no dependence* on the dimension d . Note that one can relax the assumption of $\int_{\mathcal{X}} \rho_*^4(x) \mathcal{P}_X(dx) \leq C$ to $\int_{\mathcal{X}} \rho_*^2(x) \mathcal{P}_X(dx) \leq C$ at the cost of the factor δ^{-1} instead of $\delta^{-1/2}$ in the above statement.

4. Applications to Wide Neural Networks

Before extending the results to neural networks, we need to study a Neural-Tangent-type kernel defined below in (10), which is slightly different from the inner-product kernel as in (4). Specifically, we consider kernels of the following form:

$$k(x, x') = \|x\| \|x'\| \sum_{i=0}^{\infty} \alpha_i (\cos \angle(x, x'))^i \quad (10)$$

where $\cos \angle(x, x') = x^\top x' / \|x\| \|x'\|$. Suppose that $\alpha_i \geq 0$ for all i and $\sup\{i : \alpha_i > 0\} = \infty$. Note that the above kernel reduces to the inner-product kernel when the data lies on a fixed radius sphere $\|x\| \equiv R$.

Corollary 4 (Generalization of Neural-Tangent-Type Kernels) *Consider the type of kernels defined in (10), which subsumes the Neural Tangent Kernel as a special case. Consider n i.i.d. data pairs (x_i, y_i) drawn from $\mathcal{P}_{X,Y}$, and denote the target function $f_*(x) = \mathbb{E}[Y|X = x]$. Suppose the conditions on f_* and $\mathcal{P}_{X,Y}$ specified by Theorems 2-3 are satisfied. Consider integer ι that satisfy $d^\iota \log d \lesssim n \lesssim d^{\iota+1} / \log d$. Then the result of Theorem 1 can be generalized to such kernels: with high probability, the following upper bound on the risk holds:*

$$\mathbb{E} \left[\|\hat{f} - f_*\|_{\mathcal{P}_X}^2 | \mathbf{X} \right] \leq C \cdot \left(\frac{d^\iota}{n} + \frac{n \log d}{d^{\iota+1}} \right).$$

The connection between Neural Tangent Kernel (NTK) and wide neural networks is by now well-known (Jacot et al., 2018; Du et al., 2018; Cho and Saul, 2009; Bietti and Mairal, 2019). Sufficiently wide randomly initialized (with appropriate scaling) neural networks converge to the minimum-norm interpolating solution with respect to NTK, under appropriate assumptions. This connection allows us to leverage Corollary 4 for establishing estimation and generalization guarantees for these models.

For completeness, we show that NTK for infinitely-wide neural networks is indeed of the form (10). Here we consider a one-hidden-layer neural network defined as follows:

$$f(x; W, a) := \frac{1}{\sqrt{m}} \sum_{j=1}^m a_j \sigma(w_j^\top \tilde{x}),$$

where the input x is a d -dimensional vector $W = (w_1, \dots, w_m)$ is a $(d+1) \times m$ matrix and $a = (a_1, \dots, a_m)$ is a m -dimensional vector and $\tilde{x} = (x^\top, \sqrt{d})^\top$. The NTK h^m is defined by

$$h^m(x, x') := \frac{1}{m} \left(\sum_{j=1}^m \sigma(w_j^\top \tilde{x}) \sigma(w_j^\top \tilde{x}') + x^\top x' \sum_{j=1}^m a_j^2 \sigma'(w_j^\top \tilde{x}) \sigma'(w_j^\top \tilde{x}') \right).$$

Assume that the parameters are initialized according to i.i.d. $\mathcal{N}(0, 1)$. Then the above kernel converges pointwise to the following kernel as $m \rightarrow \infty$:

$$\begin{aligned} h^\infty(x, x') &:= \mathbb{E}_{w \sim \mathcal{N}(0, I_{d+1})} \left(\sigma(w^\top \tilde{x}) \sigma(w^\top \tilde{x}') + \tilde{x}^\top \tilde{x}' \sigma'(w^\top \tilde{x}) \sigma'(w^\top \tilde{x}') \right) \\ &= \frac{1}{4\pi} \|\tilde{x}\| \|\tilde{x}'\| \left((\pi - \angle(\tilde{x}, \tilde{x}')) \cos \angle(\tilde{x}, \tilde{x}') + \sin \angle(\tilde{x}, \tilde{x}') \right) + \frac{1}{2\pi} \tilde{x}^\top \tilde{x}' (\pi - \angle(\tilde{x}, \tilde{x}')) \\ &= \frac{1}{4\pi} \|\tilde{x}\| \|\tilde{x}'\| U(\cos \angle(\tilde{x}, \tilde{x}')) \end{aligned}$$

and U takes the following analytic form

$$\begin{aligned} U(t) &:= 3t(\pi - \arccos(t)) + \sqrt{1-t^2} \\ &= 1 + \frac{3\pi t}{2} + \sum_{i=0}^{\infty} \left(\frac{3(\frac{1}{2})_i}{(1+2i)i!} - \frac{1}{2} \frac{(\frac{1}{2})_i}{(i+1)!} \right) t^{2i+2} \\ &= 1 + \frac{3\pi t}{2} + \sum_{i=0}^{\infty} \frac{(4i+5)(\frac{1}{2})_i t^{2i+2}}{2(2i+1)(i+1)i!} \end{aligned}$$

where $(\frac{1}{2})_i = \frac{1}{2} \times \frac{3}{2} \times \dots \times (\frac{1}{2} + i - 1)$ is the Pochhammer symbol. Now we have verified that the Neural Tangent Kernel h^∞ is of the form (10).

In fact, it is not difficult to prove that for multilayer fully connected neural network the NTK is also of this form with all positive Taylor coefficients if seen as a function of $\cos \angle(\tilde{x}, \tilde{x}')$.

5. Proof of Theorem 2

In this section we prove Theorem 2, as it sheds light on the emergence of the multiple descent phenomenon. We need only prove (i) because (ii) shall follow easily from the proof of (i). We will first show that with high probability, \mathbf{K} is invertible, and $n\mathbf{K} \succ c\mathbf{I}_n$ for some constant $c > 0$. To show this, it suffices to prove

$$n\mathbf{K}^{[l']} \succ c\mathbf{I}_n$$

Write $\mathbf{K}^{[l']}$ as

$$n\mathbf{K}^{[l']} = A + B$$

where A is the diagonal terms and B is the non-diagonal terms. With probability at least $1 - n^{-C}$, we shall have

$$\lambda_{\min}(A) = \Omega(1). \quad (11)$$

and by Hölder's inequality on matrix induced norm,

$$\|B\|_{\text{op}} = \|B\|_{\ell_2 \rightarrow \ell_2} \leq \sqrt{\|B\|_{\ell_1 \rightarrow \ell_1} \|B\|_{\ell_\infty \rightarrow \ell_\infty}} \quad (12)$$

$$= \|B\|_{\ell_1 \rightarrow \ell_1} \leq O \left(n \times \left(\frac{\sqrt{\log n}}{\sqrt{d}} \right)^{l'} \right) \leq O \left(\frac{1}{\sqrt{d}} \right). \quad (13)$$

Hence

$$n\mathbf{K} \succeq n\mathbf{K}^{[\ell']} \succeq (\lambda_{\min}(A) - \|B\|_{\text{op}})\mathbf{I}_n \succeq c\mathbf{I}_n . \quad (14)$$

Now back to the proof of (i). Recall the normalized kernel matrix $\mathbf{K} = k(\mathbf{X}, \mathbf{X})/n$.

$$\mathbb{E}_x \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 \quad (15)$$

$$\lesssim \sum_{i=0}^{\ell} \mathbb{E}_x \|\mathbf{K}^{-1} \frac{1}{n} (\mathbf{X}x)^i / d^i\|^2 + \mathbb{E}_x \|\mathbf{K}^{-1} \frac{1}{n} \sum_{i=\ell+1}^{\infty} (\mathbf{X}x)^i / d^i\|^2 \quad (16)$$

$$\lesssim \frac{1}{n^2} \sum_{i=0}^{\ell} \mathbb{E}_x \|\mathbf{K}^{-1} (\mathbf{X}x)^i / d^i\|^2 + \|(\mathbf{K}\mathbf{K})^{-1}\|_{\text{op}}^2 \cdot \mathbb{E}_x \left\| \sum_{i=\ell+1}^{\infty} (\mathbf{X}x)^i / d^i \right\|^2 \quad (17)$$

$$\lesssim \frac{1}{n^2} \sum_{i=0}^{\ell} \mathbb{E}_x \left[\|(\mathbf{K}^{[\leq i]})^+\|_{\text{op}}^2 \cdot \|(\mathbf{X}x)^i / d^i\|^2 \right] + \frac{n}{d^{\ell+1}} \quad (18)$$

Now, by Proposition 1, the above quantity is at most

$$\begin{aligned} &\lesssim \frac{1}{n^2} \sum_{i=0}^{\ell} \mathbb{E}_x [d^{2i} \cdot \|(\mathbf{X}x)^i / d^i\|^2] + \frac{n}{d^{\ell+1}} \\ &\lesssim \frac{1}{n^2} \sum_{i=0}^{\ell} \mathbb{E}_x \|(\mathbf{X}x)^i\|^2 + \frac{n}{d^{\ell+1}} \lesssim \frac{1}{n^2} \sum_{i=0}^{\ell} nd^i + \frac{n}{d^{\ell+1}} \leq \frac{d^{\ell}}{n} + \frac{n}{d^{\ell+1}} . \end{aligned}$$

The line (18) requires some explanation. First observe that $(\mathbf{X}x)^i$ lies in the span of $\mathbf{K}^{[\leq i]}$. Write $\mathbf{K} = \mathbf{K}^{[\leq i]} + \mathbf{K}^{[> i]} \succeq \mathbf{K}^{[\leq i]}$, we know that when restricted to the column space Φ of $\mathbf{K}^{[\leq i]}$ (via projection operator Π_{Φ}), the operator \mathbf{K} satisfy

$$\mathbf{K}|_{\Phi} = \mathbf{K}^{[\leq i]} + \Pi_{\Phi} \mathbf{K}^{[> i]} \Pi_{\Phi} \succeq \mathbf{K}^{[\leq i]}$$

Therefore for $v = (\mathbf{X}x)^i$ in the span of Φ , we have

$$\begin{aligned} \|(\mathbf{K}|_{\Phi})^{-1} v\|^2 &\leq \|(\mathbf{K}^{[\leq i]})^+\|_{\text{op}} \cdot \|(\mathbf{K}|_{\Phi})^{-1/2} v\|^2 \\ &\leq \|(\mathbf{K}^{[\leq i]})^+\|_{\text{op}} \cdot v^{\top} (\mathbf{K}^{[\leq i]})^+ v \leq \|(\mathbf{K}^{[\leq i]})^+\|_{\text{op}}^2 \cdot \|v\|^2 . \end{aligned}$$

Note that in above derivation, we use that by concentration $\mathbb{E}_x [(x^{\top} x_j)^{2\ell}] \lesssim d^{\ell}$ in high probability over \mathbf{X} . This can be seen because conditioning on \mathbf{X} , $x^{\top} x_j$ is sub-Gaussian with parameter $\|x_j\|_2$, and

$$\mathbb{E}_x \left[(x^{\top} x_j)^{2\ell} \right] \lesssim \|x_j\|_2^{2\ell} \lesssim d^{\ell}, \quad \text{w.p. } 1 - \exp(d - \log n) .$$

This concludes the proof of Theorem 2.

6. Proof of Theorem 3

In this section we provide a proof of Theorem 3. Define the following ‘‘surrogate’’ function for analyzing the bias term $\tilde{f}_n(x) := \frac{1}{n} \sum_{i=1}^n k(x, x_i) \rho_*(x_i)$. We start with splitting

$$\begin{aligned} &\|k(x, \mathbf{X}) k(\mathbf{X}, \mathbf{X})^{-1} f_*(\mathbf{X}) - f_*(x)\|_{\mathcal{P}_X}^2 \\ &\lesssim \left\| k(x, \mathbf{X}) k(\mathbf{X}, \mathbf{X})^{-1} \left[f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X}) \right] \right\|_{\mathcal{P}_X}^2 + \|k(x, \mathbf{X}) k(\mathbf{X}, \mathbf{X})^{-1} \tilde{f}_n(\mathbf{X}) - f_*(x)\|_{\mathcal{P}_X}^2 . \end{aligned}$$

The first term is equal to $\int \left\langle k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, x), f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X}) \right\rangle^2 \mathcal{P}_X(dx)$, which is at most

$$\begin{aligned} & \int \|k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, x)\|^2 \mathcal{P}_X(dx) \cdot \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 \\ & = \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 \cdot \mathbb{E}_{x \sim \mathcal{P}_X} \|k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, x)\|^2 \end{aligned}$$

by the Cauchy-Schwartz inequality. By Proposition 2 in the Appendix,

$$\mathbb{E}_{\mathbf{X}} \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 \lesssim 1.$$

For the second term, defining a vector $\tilde{V} = [\frac{\rho_*(x_1)}{n}, \dots, \frac{\rho_*(x_n)}{n}] \in \mathbb{R}^n$, we have $\tilde{f}_n(\mathbf{X}) = k(\mathbf{X}, \mathbf{X})\tilde{V}$. Then

$$\begin{aligned} \|k(x, \mathbf{X})k(\mathbf{X}, \mathbf{X})^{-1}\tilde{f}_n(\mathbf{X}) - f_*(x)\|_{\mathcal{P}_X}^2 & = \|k(x, \mathbf{X})k(\mathbf{X}, \mathbf{X})^{-1}k(\mathbf{X}, \mathbf{X})\tilde{V} - f_*(x)\|_{\mathcal{P}_X}^2 \\ & = \|\tilde{f}_n(x) - f_*(x)\|_{\mathcal{P}_X}^2. \end{aligned}$$

By Proposition 3 in the Appendix, it holds that $\mathbb{E}_{\mathbf{X}} \|f_*(x) - \tilde{f}_n(x)\|_{\mathcal{P}_X}^2 \lesssim \frac{1}{n}$.

7. Discussion

We showed that minimum norm interpolants in RKHS, under the assumptions employed in this paper, have risk that vanishes in n for a wide range of scalings $d \asymp n^\alpha$, $\alpha \in (0, 1)$. Notably, the places where our upper bounds become vacuous are fractions $\alpha = 1/i$ for integer i . The phenomenon of non-monotonicity with peaks at these locations is supported by empirical evidence, and generalizes the double-descent behavior in linear regression and other models.

A more precise description of the risk curve is an interesting research direction. In addition, it would be interesting to understand the effect of regularization on the peaks. In terms of assumptions, the i.i.d. assumption on the coordinates can certainly be lifted, and we believe similar results hold under a rotation of vectors with independent or weakly-dependent coordinates. Some degree of independence, however, is needed to capture the scaling with d .

Finally, we mention that the difficulty of analyzing min-norm interpolants is greatly reduced in the *noiseless* case when $y_i = f_*(x_i)$. Indeed, in this case the variance term is zero. Moreover, one can appeal to known results on lower isometry (for instance, Lemmas 8 and 9 in (Rakhlin et al., 2017)) to establish that, up to polylogarithmic factors,

$$\mathbb{E} \|\hat{f} - f_*\|_{L_2(\mathcal{P}_X)}^2 \lesssim \mathcal{R}_n(\mathcal{F})^2,$$

the squared Rademacher averages of \mathcal{F} , whenever $\hat{f} \in \mathcal{F}$. Since \hat{f} is the min-norm interpolant, we can take \mathcal{F} to be the ball in \mathcal{H} of radius $\|f_*\|_K$, yielding a consistency result

$$\mathbb{E} \|\hat{f} - f_*\|_{L_2(\mathcal{P}_X)}^2 \lesssim \frac{\|f_*\|_K^2}{n}.$$

This can be further tightened to an upper bound in terms of $\|\hat{f}\|_K^2$, in high probability. In contrast, the norm $\|\hat{f}\|_K$ is not easily controlled in the noisy case.

References

- Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *arXiv preprint arXiv:1906.11300*, 2019.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning and the bias-variance trade-off. *arXiv preprint arXiv:1812.11118*, 2018a.
- Mikhail Belkin, Daniel Hsu, and Partha Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. *arXiv preprint arXiv:1806.05161*, 2018b.
- Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand kernel learning. *arXiv preprint arXiv:1802.01396*, 2018c.
- Mikhail Belkin, Alexander Rakhlin, and Alexandre B Tsybakov. Does data interpolation contradict statistical optimality? *arXiv preprint arXiv:1806.09471*, 2018d.
- Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. In *Advances in Neural Information Processing Systems*, pages 12873–12884, 2019.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Youngmin Cho and Lawrence K Saul. Kernel methods for deep learning. In *Advances in neural information processing systems*, pages 342–350, 2009.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- Vitaly Feldman. Does learning require memorization? a short tale about a long tail. *arXiv preprint arXiv:1906.05271*, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *arXiv preprint arXiv:1904.12191*, 2019.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Vladimir Koltchinskii and Shahar Mendelson. Bounding the smallest singular value of a random matrix without concentration. *International Mathematics Research Notices*, 2015(23):12991–13008, 2015.
- Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel” ridgeless” regression can generalize. *The Annals of Statistics*, to appear, 2018.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.

Shahar Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39, 2014.

Alexander Rakhlin and Xiyu Zhai. Consistency of interpolation with laplace kernels is a high-dimensional phenomenon. *arXiv preprint arXiv:1812.11167*, 2018.

Alexander Rakhlin, Karthik Sridharan, and Alexandre B. Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824, May 2017. doi: 10.3150/14-bej679. URL <https://doi.org/10.3150/14-bej679>.

Appendix A. Proof of Main Proposition

The proof aims to establish the *restricted lower isometry* behavior for the empirical kernel \mathbf{K} when restricting to the eigen-space of the population covariance operator with rank $\binom{i+d}{i}$ (sorted according to the eigenvalues). We show a lower bound for the restricted lower isometry, as multiplicatively equivalent to the population eigenvalues. The approach proceeds along the lines of (Koltchinskii and Mendelson, 2015; Mendelson, 2014). One technical contribution is establishing the “small-ball” property for the polynomial basis of the kernel.

A.1. Preparation

We use the notation “ r_i ,” to represent a sequence of indices $r_1 \cdots r_d$. For example, we can use this to abbreviate a monomial with order $r_1 + \dots + r_d$: for monomial p (recall that x is a d -dimensional vector),

$$p_{r_i}(x) := p_{r_1 \cdots r_d}(x) = \prod_{i=1}^d (x[i])^{r_i} \quad (19)$$

where $x[i]$ denotes the i -th coordinate of x . This notation is also used in tensors T , for example:

$$T_{r_i} := T_{r_1 \cdots r_d}. \quad (20)$$

First, we fix an index $\iota \leq \iota_0$. After we prove it for ι , the conclusion shall follow easily from a union bound over $\iota = 1, \dots, \iota_0$.

Consider the Taylor expansion, expressed with the multi-index r_1, \dots, r_d

$$n\mathbf{K}_{ij}^{[\leq \iota]} = \sum_{r_1, \dots, r_d \geq 0}^{\iota} c_{r_1 \cdots r_d} \alpha_{r_1 + \dots + r_d} p_{r_1 \cdots r_d}(x_i) p_{r_1 \cdots r_d}(x_j) / d^{r_1 + \dots + r_d} \quad (21)$$

where

$$c_{r_1 \cdots r_d} = \frac{(r_1 + \dots + r_d)!}{r_1! \cdots r_d!}, \quad p_{r_1 \cdots r_d}(x_i) = (x_i[1])^{r_1} \cdots (x_i[d])^{r_d}. \quad (22)$$

Fix an ordering of all (r_1, \dots, r_d) such that $r_i \geq 0$, $\sum r_i \leq \iota$, and let

$$(r_1 \cdots r_d)_\iota \in \left\{ 1, 2, \dots, \binom{\iota + d}{\iota} \right\}$$

denote the index of (r_1, \dots, r_d) in the ordering. Define the $n \times \binom{\iota+d}{\iota}$ matrix Φ as

$$\Phi_{i,(r_1 \dots r_d)_\iota} = \sqrt{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}} p_{r_1 \dots r_d}(x_i) / d^{(r_1 + \dots + r_d)/2}. \quad (23)$$

Then it is not hard to see

$$\mathbf{K}^{[\leq \iota]} = \frac{1}{n} \Phi \Phi^\top, \quad (24)$$

which has the same nonzero spectrum as the covariance operator

$$\Theta = \frac{1}{n} \Phi^\top \Phi. \quad (25)$$

We know

$$\Theta_{(r_1 \dots r_d)_\iota, (r'_1 \dots r'_d)_\iota} = \frac{1}{n} \sum_{i=1}^n \sqrt{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}} \sqrt{c_{r'_1 \dots r'_d} \alpha_{r'_1 + \dots + r'_d}} \frac{p_{r_1 \dots r_d}(x_i) p_{r'_1 \dots r'_d}(x_i)}{d^{(r_1 + \dots + r_d + r'_1 + \dots + r'_d)/2}}. \quad (26)$$

It would be hard to work directly with Φ to analyze the eigenvalues of the random matrix because of complex correlation structure in the entries. Instead, we define another matrix Ψ with the following properties:

1. $\Psi^\top \Psi$ is easier to analyze from a probabilistic point of view;
2. there is a linear transformation Λ with bounded operator norm $\|\Lambda\|_{\ell_2 \rightarrow \ell_2}, \|\Lambda^{-1}\|_{\ell_2 \rightarrow \ell_2}$ such that

$$\Phi = \Psi \Lambda. \quad (27)$$

With such a Λ , we have

$$\Theta = \Lambda^\top \Psi^\top \Psi \Lambda. \quad (28)$$

Such Λ can be obtained through the Gram-Schmidt process on the polynomial basis, which we will elaborate on next.

A.2. Gram-Schmidt Process

We now proceed in the following steps:

Step 1. Define Ψ Given a distribution \mathcal{P} over \mathbb{R} , define $q_0, q_1, \dots, q_k, \dots$, to be the sequence of polynomials obtained by the Gram-Schmidt process on the basis $\{1, t, \dots, t^k, \dots\}$ w.r.t. the inner product of space $L^2(\mathcal{P})$. Define the $n \times \binom{\iota+d}{\iota}$ matrix Ψ as

$$\Psi_{i,(r_1 \dots r_d)_\iota} = \sqrt{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}} q_{r_1 \dots r_d}(x_i) / d^{(r_1 + \dots + r_d)/2}. \quad (29)$$

Here $q_{r_1 \dots r_d}$ will be defined in (32).

To be concrete, one can see that

$$q_0 = 1, q_1 = (x - m_1) / \sqrt{(m_2 - m_1^2)}, \dots, \quad (30)$$

where m_i is the i th moment of \mathcal{P} , and for $i < j$

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{P}} [\mathbf{x}^i q_j(\mathbf{x})] = 0. \quad (31)$$

The following lemmas on properties of the polynomial basis due to Gram-Schmidt process will be useful.

Lemma 5 Suppose that $f(x) = \sum_{r_i} u_{r_i}; q_{r_i};(x)$, then

$$u_{r_i}; = \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [f(X) q_{r_i};(X)], \text{ where } q_{r_i};(x) := \prod_{i=1}^d q_{r_i}(x[i]). \quad (32)$$

Lemma 6 $\mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) p(X)] = 0$ if there exists one i such that the degree of $p(x)$ as a polynomial of $x[i]$ is less than r_i .

Lemma 7 (Triangle condition) $\mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) q_{r'_i};(X) q_{r''_i};(X)] = 0$ if there exists one i , such that $r_i > r'_i + r''_i$.

Step 2. Existence of Λ .

Lemma 8 There is a $\binom{\iota+d}{\iota} \times \binom{\iota+d}{\iota}$ invertible matrix Λ such that

$$\Phi = \Psi \Lambda. \quad (33)$$

Proof [Proof of Lemma 8] For any $a \in \mathbb{R}^{\binom{\iota+d}{\iota}}$, there is a unique $b \in \mathbb{R}^{\binom{\iota+d}{\iota}}$ such that

$$\begin{aligned} & \sum_{r_1+\dots+r_d \leq \iota, r_i \geq 0} a_{(r_1 \dots r_d)\iota} \sqrt{c_{r_1 \dots r_d} \alpha_{r_1+\dots+r_d}} p_{r_1 \dots r_d} / d^{(r_1+\dots+r_d)/2} \\ &= \sum_{r_1+\dots+r_d \leq \iota, r_i \geq 0} b_{(r_1 \dots r_d)\iota} \sqrt{c_{r_1 \dots r_d} \alpha_{r_1+\dots+r_d}} q_{r_1 \dots r_d} / d^{(r_1+\dots+r_d)/2}. \end{aligned} \quad (34)$$

As a result

$$\Phi a = \Psi b.$$

Choose Λ to be the linear mapping that maps a to b

$$\Phi a = \Psi \Lambda a.$$

This holds for any a , so we have

$$\Phi = \Psi \Lambda. \quad \blacksquare$$

Step 3. Boundedness of Λ and Λ^{-1} . For a vector $v \in \mathbb{R}^{\binom{\iota+d}{\iota}}$, let $v_{\geq \iota'}$ be the vector such that

$$(v_{\geq \iota'})_{(r_1 \dots r_d)\iota} = v_{(r_1 \dots r_d)\iota} \mathbb{1}_{\{r_1+\dots+r_d \geq \iota'\}}. \quad (35)$$

Define similarly for $v_{\iota'}, v_{> \iota'}, \Lambda_{\geq \iota'}, \Lambda_{\geq \iota', < \iota'}$.

Lemma 9 There is a constant $C(\iota)$ independent of d such that

$$\|\Lambda\|_{\ell_2 \rightarrow \ell_2}, \|\Lambda^{-1}\|_{\ell_2 \rightarrow \ell_2} \leq C(\iota). \quad (36)$$

Proof [Proof of Lemma 9] We start with few claims about the Gram-Schmidt process and the structure of Λ .

Claim 1: $\Lambda_{\geq \iota', < \iota'} = 0$ for any ι' . Alternatively, if $b = \Lambda a$ then

$$b_{\geq \iota'} = \Lambda_{\geq \iota'} a_{\geq \iota'}. \quad (37)$$

Proof [Proof of Claim 1]

We need only to show that if $a_{\geq \iota'} = 0$, then $b_{\geq \iota'} = 0$. Observe that $a_{\geq \iota'} = 0$ implies that the left hand of equation (34) is of degree less than ι' . Since this is an equality, the right hand side must be of degree less than ι' . Note that this implies that $b_{\geq \iota'} = 0$. \blacksquare

Claim 2: $\Lambda_{\iota', \iota'}$ is diagonal with

$$c \leq \lambda_{\min}(\Lambda_{\iota', \iota'}) \leq \lambda_{\max}(\Lambda_{\iota', \iota'}) \leq C \quad (38)$$

where c, C depend only on ι', \mathcal{P} .

Proof [Proof of Claim 2]

Given r_i ; and r'_i ; with $\sum_i r_i = \sum_i r'_i = \iota'$, we have

$$\Lambda_{(r_i)_{\iota'}, (r'_i)_{\iota'}} = \frac{\sqrt{c_{r'_1 \dots r'_d}}}{\sqrt{c_{r_1 \dots r_d}}} \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i}(X) p_{r'_i}(X)]. \quad (39)$$

If $r_i \neq r'_i$; there is at least one i such that $r_i > r'_i$, then according to Lemma 6, we have

$$\Lambda_{(r_1 \dots r_d)_{\iota'}, (r'_1 \dots r'_d)_{\iota'}} = 0 \quad (40)$$

Therefore $\Lambda_{\iota', \iota'}$ is diagonal. Now we have

$$\begin{aligned} \Lambda_{(r_1 \dots r_d)_{\iota'}, (r_1 \dots r_d)_{\iota'}} &= \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i}(X) p_{r_i}(X)] \\ &= \prod_{i=1}^d \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} [q_{r_i}(\mathbf{x}) p_{r_i}(\mathbf{x})]. \end{aligned} \quad (41)$$

Note that $q_0 \equiv 1$. Since the set $I := \{1 \leq i \leq d : r_i \text{ is nonzero}\}$ is of size at most ι , $\Lambda_{(r_i)_{\iota'}, (r_i)_{\iota'}}$ is uniformly bounded by the following constant (depending on ι and \mathcal{P}):

$$\left(\max_{r_i} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} [q_{r_i}(\mathbf{x}) p_{r_i}(\mathbf{x})] \right)^{\iota}. \quad (42)$$

Claim 3: Let $b_{\iota'} = \Lambda_{\iota', \iota'} a_{\iota'} + \Lambda_{\iota', > \iota'} a_{> \iota'}$. Then $\|\Lambda_{\iota', > \iota'}\|_{\ell_2 \rightarrow \ell_2}$ has an upper bound that only depends on ι .

Proof [Proof of Claim 3]

An entry $(\Lambda_{\iota', > \iota'})_{(r_1 \dots r_d)_{\iota'}, (r'_1 \dots r'_d)_{\iota'}}$ for $\sum_i r_i = \iota'$ can be obtained by

$$\begin{aligned} &(\Lambda_{\iota', > \iota'})_{(r_1 \dots r_d)_{\iota'}, (r'_1 \dots r'_d)_{\iota'}} \\ &= \sqrt{\frac{c_{r'_1 \dots r'_d} \alpha_{r'_1 + \dots + r'_d}}{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}}} d^{(r_1 + \dots + r_d - r'_1 - \dots - r'_d)/2} \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i}(X) p_{r'_i}(X)] \\ &= \sqrt{\frac{c_{r'_1 \dots r'_d} \alpha_{r'_1 + \dots + r'_d}}{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}}} d^{(r_1 + \dots + r_d - r'_1 - \dots - r'_d)/2} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} [q_{r_i}(\mathbf{x}) p_{r'_i}(\mathbf{x})]. \end{aligned}$$

Only when $\forall i, r_i \leq r'_i$, the above term is nonzero and scales with d in the order of $d^{(r_1+\dots+r_d-r'_1-\dots-r'_d)/2}$. As a result,

$$\begin{aligned}
 b_{(r_1 \dots r_d)\iota}^2 &= \left(\sum_{r'_1, \dots, r'_d} (\Lambda_{\iota', > \iota'})(r_1 \dots r_d)\iota, (r'_1 \dots r'_d)\iota a_{(r'_1 \dots r'_d)\iota} \right)^2 \\
 &= \left(\sum_{r_i \leq r'_i \forall i, \sum_i r'_i > \iota'} (\Lambda_{\iota', > \iota'})(r_1 \dots r_d)\iota, (r'_1 \dots r'_d)\iota a_{(r'_1 \dots r'_d)\iota} \right)^2 \\
 &\leq \left(\sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} (\Lambda_{\iota', > \iota'})(r_1 \dots r_d)\iota, (r'_1 \dots r'_d)\iota \right) \left(\sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2 \right) \\
 &\leq \left(\sum_{l \geq 1} \sum_{\substack{r_i \leq r'_i \\ \sum_i r'_i = \iota' + l}} (\Lambda_{\iota', > \iota'})(r_1 \dots r_d)\iota, (r'_1 \dots r'_d)\iota \right) \left(\sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2 \right) \\
 &\lesssim \left(\sum_{l \geq 1} \sum_{\substack{r_i \leq r'_i \\ \sum_i r'_i = \iota' + l}} (d^{-l/2})^2 \right) \left(\sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2 \right) \\
 &\lesssim \left(\sum_{l \geq 1} d^l \times (d^{-l/2})^2 \right) \left(\sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2 \right) \\
 &\lesssim \sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2.
 \end{aligned}$$

Then we have

$$\begin{aligned}
 &\sum_{\sum_i r_i = \iota'} b_{(r_1 \dots r_d)\iota}^2 \\
 &\lesssim \sum_{\sum_i r_i = \iota'} \sum_{r_i \leq r'_i, \sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2 \\
 &\lesssim \sum_{\sum_i r'_i > \iota'} a_{(r'_1 \dots r'_d)\iota}^2
 \end{aligned}$$

where the last inequality holds because for a fixed r'_1, \dots, r'_d with $r'_1 + \dots + r'_d = \iota$, there is at most 2^ι choices of r_1, \dots, r_d such that $\forall i, 0 \leq r_i \leq r'_i$. \blacksquare

Using induction on $\iota' \leq \iota$ backwards will complete the proof that $\|b\| \asymp \|a\|$. Specifically, assume that $\|b_{> \iota'}\| \asymp \|a_{> \iota'}\|$, then since

$$\|b_{\iota'}\| \leq \|\Lambda_{\iota', \iota'} a_{\iota'}\| + \|\Lambda_{\iota', > \iota'} a_{> \iota'}\| \lesssim \|a_{\geq \iota'}\|, \quad (43)$$

we have

$$\|b_{\geq \iota'}\| \leq \|b_{\iota'}\| + \|b_{> \iota'}\| \lesssim \|a_{\geq \iota'}\|. \quad (44)$$

For the other direction, we have

$$\|a_{\iota'}\| = \|\Lambda_{\iota'}^{-1}(b_{\iota'} - \Lambda_{\iota', > \iota'} a_{> \iota'})\| \lesssim \|b_{\iota'} - \Lambda_{\iota', > \iota'} a_{> \iota'}\| \lesssim \|b_{\iota'}\| + \|a_{> \iota'}\| \lesssim \|b_{\iota'}\| + \|b_{> \iota'}\| \lesssim \|b_{\geq \iota'}\|.$$

Then

$$\|a_{\geq \iota'}\| \lesssim \|b_{\geq \iota'}\|. \quad (45)$$

■

Now we can proceed to analyze the spectrum of $\Psi^\top \Psi$, given the boundedness of Λ .

A.3. Small Ball Property

Define the following function over \mathbb{R}^d indexed by $u \in \mathbb{R}^{\binom{\iota+d}{\iota}}$

$$f_u(x) = \sum_{r_1, \dots, r_d \geq 0, \sum_i r_i \leq \iota} u_{(r_1 \dots r_d)_\iota} \sqrt{c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d}} q_{r_1 \dots r_d}(x) / d^{(r_1 + \dots + r_d)/2}. \quad (46)$$

In this section, we will prove that there exist constants $0 < \epsilon, \delta < 1$, such that for any u with $\|u\| = 1$,

$$\mathbb{P}(f_u(X)^2 > \epsilon \mathbb{E}[f_u(X)^2]) \geq \delta. \quad (47)$$

This is so called the *small-ball property* for the random variable $f_u(X)$, with $X \sim \mathcal{P}^{\otimes d}$.

Claim 1: $\forall u \in \mathbb{R}^{\binom{\iota+d}{\iota}}$ with $\|u\|_2 = 1$, there are constants $\delta, \epsilon > 0$ depending only on ι, \mathcal{P} such that with probability at least $1 - \delta$

$$f_u(x_i)^2 > \epsilon d^{-\iota}. \quad (48)$$

Proof of Claim 1. First, according to the Paley-Zygmund inequality for $X \sim \mathcal{P}^{\otimes d}$ and any $0 \leq \theta \leq 1$,

$$\mathbb{P}(f_u(X)^2 > \theta \mathbb{E}[f_u(X)^2]) \geq (1 - \theta)^2 \frac{\mathbb{E}[f_u(X)^2]^2}{\mathbb{E}[f_u(X)^4]}. \quad (49)$$

Therefore we just need to show that

$$\mathbb{E}[f_u(X)^2] \gtrsim d^{-\iota}, \quad (50)$$

and

$$\mathbb{E}[f_u(X)^4] \lesssim \mathbb{E}[f_u(X)^2]^2. \quad (51)$$

For equation (50), we have by the orthogonality of q_{r_i} ,

$$\begin{aligned} \mathbb{E}[f_u(X)^2] &= \sum_{r_1, \dots, r_d \geq 0, \sum_i r_i \leq \iota} u_{(r_1 \dots r_d)_\iota}^2 c_{r_1 \dots r_d} \alpha_{r_1 + \dots + r_d} / d^{r_1 + \dots + r_d} \\ &\gtrsim \sum_{r_1, \dots, r_d \geq 0, \sum_i r_i \leq \iota} u_{(r_1 \dots r_d)_\iota}^2 / d^{r_1 + \dots + r_d} \gtrsim d^{-\iota}. \end{aligned} \quad (52)$$

Equation (51) is more technical to establish, which we prove through the following lemma.

Lemma 10 Let $f_\gamma(x) := \sum_{r_1, \dots, r_d \geq 0, \sum_i r_i \leq \iota} \gamma_{r_1 \dots r_d} q_{r_1 \dots r_d}(x)$, then

$$\mathbb{E} f_\gamma(X)^4 \lesssim (\mathbb{E} f_\gamma(X))^2 = \left(\sum_{\sum_i r_i \leq \iota} \gamma_{r_1 \dots r_d}^2 \right)^2. \quad (53)$$

Proof [Proof of Lemma 10.] Write

$$f_\gamma(x)^2 = \sum_{\sum_i r_i \leq 2\iota} \theta_{r_1 \dots r_d} q_{r_1 \dots r_d}(x). \quad (54)$$

Since

$$f_\gamma(x)^2 = \sum_{r'_i; r''_i} \gamma_{r'_i; \gamma_{r''_i}; q_{r'_i};(x) q_{r''_i};(x), \quad (55)$$

then by triangle condition there are coefficients $T_{r'_i; r''_i}^{r_i};$, $M_{r'_i; r''_i}^{r_i};$ (defined in (56) and (57)) such that

$$\begin{aligned} \theta_{r_1 \dots r_d} &= \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) f_\gamma(X)^2] \\ &= \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} \left\{ q_{r_i};(X) \sum_{r'_i; r''_i} \gamma_{r'_i; \gamma_{r''_i}; q_{r'_i};(X) q_{r''_i};(X) \right\} \\ &= \sum_{r'_i; r''_i} \gamma_{r'_i; \gamma_{r''_i};} \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) q_{r'_i};(X) q_{r''_i};(X)] \\ &= \sum_{\forall i, r'_i + r''_i \geq r_i} T_{r'_i; r''_i}^{r_i}; \gamma_{r'_i}; \gamma_{r''_i}; \\ &= \sum_{s'_i + s''_i = r_i} \sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i; r''_i}^{r_i}; \neq 0}} T_{r'_i; r''_i}^{r_i}; \gamma_{r'_i}; \gamma_{r''_i}; / M_{r'_i; r''_i}^{r_i}; \end{aligned}$$

where the coefficients are given by

$$T_{r'_i; r''_i}^{r_i}; := \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) q_{r'_i};(X) q_{r''_i};(X)], \quad (56)$$

and

$$M_{r'_i; r''_i}^{r_i}; := \#\{(s'_i; s''_i) : s'_i + s''_i = r_i, s'_i \leq r'_i, s''_i \leq r''_i\}. \quad (57)$$

Now we will upper bound $T_{r'_i; r''_i}^{r_i};$

$$\begin{aligned} &|T_{r'_i; r''_i}^{r_i};| \\ &= \left| \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} [q_{r_i};(X) q_{r'_i};(X) q_{r''_i};(X)] \right| \\ &= \left| \mathbb{E}_{X \sim \mathcal{P}^{\otimes d}} \prod_{i=1}^d q_{r_i}(X[i]) q_{r'_i}(X[i]) q_{r''_i}(X[i]) \right| \\ &= \left| \prod_{i=1}^d \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} q_{r_i}(\mathbf{x}) q_{r'_i}(\mathbf{x}) q_{r''_i}(\mathbf{x}) \right|. \end{aligned}$$

Note that $q_0 \equiv 1$. Since the set $I := \{i \in [d] : \text{one of } r_i, r'_i, r''_i \text{ is nonzero}\}$ is of size at most 3ι , $T_{r'_i, r''_i}^{r_i}$ is uniformly bounded by the following constant (depending on ι and \mathcal{P}) for $\sum r_i, \sum r'_i, \sum r''_i \leq \iota$:

$$\left(\max_{r_i, r'_i, r''_i} \mathbb{E}_{X \sim \mathcal{P}} \left[q_{r_i}(X) q_{r'_i}(X) q_{r''_i}(X) \right] \right)^{3\iota}. \quad (58)$$

As a result, we have

$$\begin{aligned} & \theta_{r_1 \dots r_d}^2 \\ &= \left(\sum_{s'_i + s''_i = r_i} \sum_{s'_i \leq r'_i, s''_i \leq r''_i} T_{r'_i, r''_i}^{r_i}; \gamma_{r'_i}; \gamma_{r''_i}; / M_{r'_i, r''_i}^{r_i} \right)^2 \\ &\lesssim \sum_{s'_i + s''_i = r_i} \left(\sum_{s'_i \leq r'_i, s''_i \leq r''_i} T_{r'_i, r''_i}^{r_i}; \gamma_{r'_i}; \gamma_{r''_i}; / M_{r'_i, r''_i}^{r_i} \right)^2 \quad \text{ignoring constants depending only on } \iota \\ &\lesssim \sum_{s'_i + s''_i = r_i} \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i, r''_i}^{r_i}; \neq 0}} (T_{r'_i, r''_i}^{r_i}; / M_{r'_i, r''_i}^{r_i})^2 \gamma_{r'_i}^2 \right) \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i, r''_i}^{r_i}; \neq 0}} \gamma_{r''_i}^2 \right) \quad \text{Cauchy inequality} \\ &\lesssim \sum_{s'_i + s''_i = r_i} \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i, r''_i}^{r_i}; \neq 0}} \gamma_{r'_i}^2 \right) \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i, r''_i}^{r_i}; \neq 0}} \gamma_{r''_i}^2 \right). \end{aligned} \quad (59)$$

Note that for $T_{r'_i, r''_i}^{r_i}; \neq 0$, we must have by the triangle condition

$$\forall i, r''_i \leq r'_i + r_i \quad (60)$$

which means that for any $r''_i \neq 0$, either $r'_i \neq 0$ or $r_i \neq 0$. Then

$$\{i \in [d] : r''_i \neq 0\} \subset \{i \in [d] : r'_i \neq 0\} \cup \{i \in [d] : r_i \neq 0\}. \quad (61)$$

As a result, for fixed r'_i ; and r_i ; , there is less than constantly many r''_i ; such that $T_{r'_i, r''_i}^{r_i}; \neq 0$. Similarly, for fixed r''_i ; and r_i ; , there is less than constantly many r'_i ; such that $T_{r'_i, r''_i}^{r_i}; \neq 0$.

We now have

$$\begin{aligned}
 \theta_{r_1 \dots r_d}^2 &\lesssim \sum_{s'_i + s''_i = r_i} \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i; r''_i}^{r_i} \neq 0}} \gamma_{r'_i}^2 \right) \left(\sum_{\substack{s'_i \leq r'_i, s''_i \leq r''_i \\ T_{r'_i; r''_i}^{r_i} \neq 0}} \gamma_{r''_i}^2 \right) \\
 &\lesssim \sum_{s'_i + s''_i = r_i} \left(\sum_{s'_i \leq r'_i} \gamma_{r'_i}^2 \right) \left(\sum_{s''_i \leq r''_i} \gamma_{r''_i}^2 \right) \\
 &= \sum_{s'_i + s''_i = r_i} \sum_{s'_i \leq r'_i} \sum_{s''_i \leq r''_i} \gamma_{r'_i}^2 \gamma_{r''_i}^2
 \end{aligned} \tag{62}$$

As a result,

$$\sum_{r_i} \theta_{r_1 \dots r_d}^2 \lesssim \sum_{r_i} \sum_{s'_i + s''_i = r_i} \sum_{s'_i \leq r'_i} \sum_{s''_i \leq r''_i} \gamma_{r'_i}^2 \gamma_{r''_i}^2 \tag{63}$$

Note that on the RHS, for a fixed r'_i, r''_i ; the term $\gamma_{r'_i}^2 \gamma_{r''_i}^2$ appears constantly many times (with constant relying on ι only), since $M_{r'_i; r''_i}^{r_i} \leq 2^{2\iota}$ and there are at most $2^{2\iota}$ number of r_i 's such that $M_{r'_i; r''_i}^{r_i} > 0$. Finally we have

$$\mathbb{E}[f_\gamma(X)^4] = \sum_{r_i} \theta_{r_1 \dots r_d}^2 \lesssim \sum_{r'_i, r''_i} \gamma_{r'_i}^2 \gamma_{r''_i}^2 = (\mathbb{E}[f_\gamma(X)^2])^2. \tag{64}$$

■

A.4. Lower Isometry

We now proceed to lower bound the smallest eigenvalue for $\frac{1}{n} \Psi^\top \Psi$, based on the small-ball property established.

Lemma 11 *With probability at least $1 - e^{-\Omega(n/d^\iota)}$, the smallest eigenvalue of $\frac{1}{n} \Psi^\top \Psi$ is larger than $Cd^{-\iota}$.*

Step 1. We will first prove: there is $\epsilon > 0$ such that for any $u \in \mathbb{R}^{\binom{\iota+d}{\iota}}$ with $\|u\|_2 = 1$,

$$\mathbb{P} \left(u^\top \Psi^\top \Psi u \geq \epsilon n d^{-\iota} \right) \geq 1 - e^{-c'n}. \tag{65}$$

Since

$$u^\top \Psi^\top \Psi u = \sum_{i=1}^n f_u(x_i)^2 \tag{66}$$

with $f_u(x_i)^2 \geq 0$ i.i.d. drawn. Define $Z_i = 1_{\{f_u(x_i)^2 \geq \epsilon d^{-\iota}\}}$. According to the Claim 1 (the small ball property (47)), we can choose ϵ such that $\mathbb{E}Z_i > C(\iota, \mathcal{P}) > 0$. Denote this $C(\iota, \mathcal{P})$ as δ . Now we have

$$u^\top \Psi^\top \Psi u \geq \epsilon d^{-\iota} \sum_{i=1}^n Z_i. \tag{67}$$

Using the Hoeffding's inequality,

$$\mathbb{P}\left(\sum_{i=1}^n Z_i \leq n(\delta - t)\right) \leq e^{-2t^2n}.$$

Take $t = \delta/2$, we get that

$$\mathbb{P}\left(u^\top \left(\frac{1}{n} \Psi^\top \Psi\right) u \geq \frac{\epsilon \delta}{2d^t}\right) \geq 1 - e^{-\delta^2 n/2}.$$

Step 2. Now we show that there is constant C such that for any $L > 0$, $u^\top (1/n \cdot \Psi^\top \Psi) u$ is L -Lipschitz w.r.t. u on the sphere, with probability at least $1 - C \left(\frac{n^2}{d} L^{-\frac{2}{\epsilon}} + n^{1-\frac{1}{2\epsilon}} L^{-\frac{1}{\epsilon}}\right)$. In fact, for $\|u\| = \|v\| = 1$, we have

$$\begin{aligned} & \|u^\top \Psi^\top \Psi u - v^\top \Psi^\top \Psi v\| \\ &= \|u^\top \Psi^\top \Psi u - v^\top \Psi^\top \Psi u + u^\top \Psi^\top \Psi v - v^\top \Psi^\top \Psi v\| \\ &= \|(u - v)^\top \Psi^\top \Psi u\| + \|(u - v)^\top \Psi^\top \Psi v\| \\ &\leq \|u - v\| \|\Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2} \|u\| + \|u - v\| \|\Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2} \|v\| \\ &= 2\|u - v\| \|\Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2}. \end{aligned}$$

Therefore, the map $u \rightarrow u^\top (1/n \cdot \Psi^\top \Psi) u$ is $2\|1/n \cdot \Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2}$ -Lipschitz. Now we need to bound the spectral norm of $\Psi^\top \Psi$. We have

$$\begin{aligned} \|1/n \cdot \Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2}^2 &= \|1/n \cdot \Lambda^{-1\top} \Phi^\top \Phi \Lambda^{-1}\|_{\ell_2 \rightarrow \ell_2}^2 \\ &\lesssim \|1/n \cdot \Phi^\top \Phi\|_{\ell_2 \rightarrow \ell_2}^2 = \|1/n \cdot \Phi \Phi^\top\|_{\ell_2 \rightarrow \ell_2}^2 \\ &= \|\mathbf{K}^{[\leq \iota]}\|_{\ell_2 \rightarrow \ell_2}^2 = \frac{1}{n^2} \sum_{1 \leq i, j \leq n} \left(h^{[\leq \iota]}(x_i^\top x_j/d)\right)^2 \end{aligned}$$

The last quantity is at most

$$\begin{aligned} & \frac{1}{n^2} \sum_{1 \leq i, j \leq n} \left(\sum_{i=0}^{\iota} \alpha_i |x_i^\top x_j/d|^i\right)^2 \lesssim \frac{1}{n^2} \sum_{i, j} \left(1 + (x_i^\top x_j/d)^{2\iota}\right) \\ &= \frac{1}{n^2} \sum_{ij} 1 + Y_{ij}^{2\iota} \end{aligned}$$

with $Y_{ij} := \frac{x_i^\top x_j}{d}$. We know that by Chebyshev's and Markov's inequality, for any $i \neq j \in [n]$, due to $\mathbb{E}[x_i[k]] = 0$,

$$\mathbb{P}\left(\frac{1}{d} \sum_{k=1}^d x_i[k] x_j[k] \geq t\right) \leq \frac{C}{dt^2}, \text{ and} \quad (68)$$

$$\mathbb{P}\left(\frac{1}{d} \sum_{k=1}^d (x_i[k])^2 \geq s\right) \leq \frac{C}{s}. \quad (69)$$

Choose $t \asymp L^{\frac{1}{i}}$, and $s \asymp n^{\frac{1}{2i}} L^{\frac{1}{i}}$. Therefore, with probability

$$1 - \frac{Cn^2}{dt^2} - \frac{Cn}{s} = 1 - C \left(\frac{n^2}{d} L^{-\frac{2}{i}} + n^{1-\frac{1}{2i}} L^{-\frac{1}{i}} \right),$$

the following bound on the Lipchitz constant holds

$$\|1/n \cdot \Psi^\top \Psi\|_{\ell_2 \rightarrow \ell_2} \lesssim 2(1 + L^2)^{1/2} \asymp L. \quad (70)$$

Step 3. Suppose $\frac{d^i \log d}{n} = o(1)$, we will show that with probability at least $1 - e^{-\Omega(n/d^i)}$,

$$\inf_{\|u\|=1} u^\top \left(\frac{1}{n} \Psi^\top \Psi \right) u = \Omega(d^{-i}). \quad (71)$$

Make an r -covering net \mathcal{N}_r of the unit sphere $\|u\| = 1$ with radius

$$r = \frac{1}{2L} \frac{\epsilon \delta}{8} d^{-i}. \quad (72)$$

Clearly, the cardinality of such r -covering is bounded by $|\mathcal{N}_r| \leq (1 + 1/\epsilon)^{\binom{i+d}{i}}$. Therefore, with probability at least $1 - (1 + 1/r)^{\binom{i+d}{i}} e^{-\delta^2 n/2}$, we know for any elements $v \in \mathcal{N}_r$ in the ϵ -cover,

$$v^\top \left(\frac{1}{n} \Psi^\top \Psi \right) v \geq \frac{\epsilon \delta}{2d^i}. \quad (73)$$

Recall that with probability at least $1 - CndL^{-\frac{2}{i}}$, $u^\top (1/n \cdot \Psi^\top \Psi)u$ is L -Lipschitz w.r.t. u on the sphere. For any u , there exists $v \in \mathcal{N}_r$, such that

$$\begin{aligned} u^\top \left(\frac{1}{n} \Psi^\top \Psi \right) u - v^\top \left(\frac{1}{n} \Psi^\top \Psi \right) v &\geq -L\|u - v\| \geq -Lr = -\frac{\epsilon \delta}{16} d^{-i} \\ u^\top \left(\frac{1}{n} \Psi^\top \Psi \right) u &\geq v^\top \left(\frac{1}{n} \Psi^\top \Psi \right) v - \frac{\epsilon \delta}{16} d^{-i} \geq \frac{\epsilon \delta}{4} d^{-i}. \end{aligned}$$

So far we have proved with that probability at least

$$1 - C \left(\frac{n^2}{d} L^{-\frac{2}{i}} + n^{1-\frac{1}{2i}} L^{-\frac{1}{i}} \right) - (1 + 1/r)^{\binom{i+d}{i}} e^{-\frac{\delta^2}{2}n} \quad (74)$$

the following holds

$$\inf_{\|u\|=1} u^\top \left(\frac{1}{n} \Psi^\top \Psi \right) u = \Omega(d^{-i}). \quad (75)$$

Now let's control the probability via a proper choice of L and r . If we choose $L = \exp(i(n - d^i \log d))n^{i-\frac{1}{2}}$, then it is easy to verify that $\frac{n^2}{d} L^{-\frac{2}{i}} + n^{1-\frac{1}{2i}} L^{-\frac{1}{i}} \leq \exp(-c'(n - d^i \log d))$, and

$$(1 + 3/r)^{\binom{i+d}{i}} e^{-\frac{\delta^2}{2}n} \leq \exp(d^i \log(d^i L/c') - cn) \leq \exp(-c''n/d^i). \quad (76)$$

Step 4. Put things together, we have shown that w.h.p.,

$$\Theta = \Phi^\top \Phi/n = \Lambda^\top \Psi^\top \Psi \Lambda/n \succ cd^{-i}. \quad (77)$$

As a result $\mathbf{K}^{[i]} = \Phi \Phi^\top/n$ has $\binom{i+d}{i}$ number of nonzero eigenvalues, all of them larger than cd^{-i} .

Appendix B. Technical Proofs

B.1. Proofs in Section 3

Proposition 2 (Leave-one-out)

$$\mathbb{E}_{\mathbf{X}} \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 \lesssim 1 \quad (78)$$

Proof [Proof of Proposition 2]

We claim that,

$$\begin{aligned} & |f_*(x_j) - \tilde{f}_n(x_j)|^2 \\ & \leq \frac{1}{n^2} |f_*(x_j) - k(x_j, x_j)\rho_*(x_j)|^2 + \frac{(n-1)^2}{n^2} \left| \frac{1}{n-1} \sum_{i \neq j} k(x_j, x_i)\rho_*(x_i) - f_*(x_j) \right|^2 \end{aligned}$$

We know that for any $x, z \in \mathcal{X}$ $k(z, x) \leq \sqrt{k(z, z)k(x, x)} \leq C$. Therefore we have

$$\mathbb{E}_{\mathbf{X}} \frac{1}{n^2} |f_*(x_j) - k(x_j, x_j)\rho_*(x_j)|^2 \leq \frac{1}{n^2},$$

since $\sup_{x \in \mathcal{X}} K(x, x) \leq C$, $\|\rho_*\|_{\mathcal{P}_X} \leq C$, and by Jensen's inequality

$$\int f_*(x)^2 \mathcal{P}_X(dx) \leq \int \int k(z, x)^2 \rho_*(z)^2 \mathcal{P}_X(dz) \mathcal{P}_X(dx) \leq C^2.$$

For the leave-on-out term,

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} \frac{(n-1)^2}{n^2} \left| \frac{1}{n-1} \sum_{i \neq j} k(x_j, x_i)\rho_*(x_i) - f_*(x_j) \right|^2 \\ & \lesssim \frac{(n-1)^2}{n^2} \mathbb{E}_{x_j} \left[\mathbb{E}_{\mathbf{X} \setminus x_j} \left| \frac{1}{n-1} \sum_{i \neq j} k(x_j, x_i)\rho_*(x_i) - f_*(x_j) \right|^2 \middle| x_j \right] \\ & \lesssim \frac{(n-1)^2}{n^2} \mathbb{E}_{x_j} \left[\frac{1}{n-1} \int k^2(x_j, x)\rho_*^2(x) \mathcal{P}_X(dx) \right] \lesssim \frac{1}{n}. \end{aligned}$$

Therefore we have

$$\mathbb{E}_{\mathbf{X}} \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 = \mathbb{E}_{\mathbf{X}} \sum_{j=1}^n |f_*(x_j) - \tilde{f}_n(x_j)|^2 \leq n \left(\frac{1}{n^2} + \frac{1}{n} \right) \lesssim 1. \quad \blacksquare$$

Proposition 3 (Variance)

$$n \mathbb{E}_{\mathbf{X}} \|f_*(x) - \tilde{f}_n(x)\|_{\mathcal{P}_X}^2 \lesssim 1 \quad (79)$$

Proof [Proof of Proposition 3]

$$\mathbb{E}_{\mathbf{X}} \int \left(f_*(x) - \tilde{f}_n(x) \right)^2 \mathcal{P}_X(dx) \leq \frac{1}{n} \int \int k^2(x, x') \rho_*^2(x') \mathcal{P}_X(dx') \mathcal{P}_X(dx) \leq C, \quad (80)$$

by boundedness of $k^2(x, x') \leq C^2$ and $\|\rho_*\|_{\mathcal{P}_X} \leq C$. ■

Proposition 4 (Probability bound) *The following bounds hold simultaneously with probability at least $1 - \delta$ on \mathbf{X} ,*

$$C_1(\mathbf{X}) := n \|f_* - \tilde{f}_n\|_{\mathcal{P}_X}^2 \lesssim \frac{1}{\sqrt{\delta}},$$

$$C_2(\mathbf{X}) := \|f_*(\mathbf{X}) - \tilde{f}_n(\mathbf{X})\|^2 = \sum_{j=1}^n \left(f_*(x_j) - \tilde{f}_n(x_j) \right)^2 \lesssim \frac{1}{\sqrt{\delta}}$$

Proof [Proof of Proposition 4] We have shown that $\mathbb{E}_{\mathbf{X}} C_1(\mathbf{X}) \lesssim 1/n$, and $\mathbb{E}_{\mathbf{X}} C_2(\mathbf{X}) \lesssim 1$. Let's use the second moment method to show the in probability bounds for both terms. Define $\tilde{h}(x, x_i) := k(x, x_i) \rho_*(x_i) - f_*(x)$. It is clear that $\mathbb{E}_{x_i \sim \mathcal{P}_X} [\tilde{h}(x, x_i)] = 0$ for any fixed x .

Second moment calculations on $C_1(\mathbf{X})$.

$$\begin{aligned} \mathbb{E} [C_1(\mathbf{X})^2] &= n^2 \mathbb{E} \left[\frac{1}{n^2} \sum_{i,j} \int (k(x, x_i) \rho_*(x_i) - f_*(x)) (k(x, x_j) \rho_*(x_j) - f_*(x)) \mathcal{P}_X(dx) \right]^2 \\ &= \frac{1}{n^2} \sum_{i,j,k,l} \mathbb{E} \left[\int \tilde{h}(x, x_i) \tilde{h}(x, x_j) \mathcal{P}_X(dx) \right] \left[\int \tilde{h}(x, x_k) \tilde{h}(x, x_l) \mathcal{P}_X(dx) \right] \end{aligned}$$

Clearly the only nonzero terms on the RHS must be either (1) $(i, j) = (k, l)$, or (2) $(i, j) \neq (k, l)$ but $i = j$ and $k = l$. In case (1), we know

$$\sum_{i,j} \mathbb{E} \left[\int \tilde{h}(x, x_i) \tilde{h}(x, x_j) \mathcal{P}_X(dx) \right]^2 \lesssim n^2.$$

In case (2), we know

$$\sum_{i \neq k} \mathbb{E} \left[\int \tilde{h}(x, x_i)^2 \mathcal{P}_X(dx) \right] \left[\int \tilde{h}(x, x_k)^2 \mathcal{P}_X(dx) \right] \lesssim n^2.$$

All other terms, must have form $\mathbb{E} \left[\int \tilde{h}(x, x_i) \tilde{h}(x, x_j) \mathcal{P}_X(dx) \right] \left[\int \tilde{h}(x, x_k)^2 \mathcal{P}_X(dx) \right] = 0$ for $i \neq j$, or $\mathbb{E} \left[\int \tilde{h}(x, x_i) \tilde{h}(x, x_j) \mathcal{P}_X(dx) \right] \left[\int \tilde{h}(x, x_k) \tilde{h}(x, x_l) \mathcal{P}_X(dx) \right] = 0$ for $i \neq j, k \neq l, (i, j) \neq (k, l)$. Therefore, we have the second moment bound $\mathbb{E} [C_1(\mathbf{X})^2] \lesssim 1$, by Chebyshev's inequality, we have the desired bound.

Second moment calculations on $C_2(\mathbf{X})$.

$$\begin{aligned}\mathbb{E} [C_2(\mathbf{X})^2] &= \sum_{i,j} \mathbb{E} \left[\left(f_*(x_i) - \tilde{f}_n(x_i) \right)^2 \left(f_*(x_j) - \tilde{f}_n(x_j) \right)^2 \right] \\ &\leq \sum_{i,j} \left[\mathbb{E} \left(f_*(x_i) - \tilde{f}_n(x_i) \right)^4 \right]^{1/2} \left[\mathbb{E} \left(f_*(x_j) - \tilde{f}_n(x_j) \right)^4 \right]^{1/2}\end{aligned}$$

We know that

$$\mathbb{E} \left(f_*(x_i) - \tilde{f}_n(x_i) \right)^4 = \frac{1}{n^4} \sum_{j_1, j_2, j_3, j_4} \mathbb{E} [\tilde{h}(x_i, x_{j_1}) \tilde{h}(x_i, x_{j_2}) \tilde{h}(x_i, x_{j_3}) \tilde{h}(x_i, x_{j_4})].$$

Divide into two cases: (1) some j equals i , (2) all j 's do not equal i . In the first case, the only nonzero terms are, of the form $\tilde{h}(x_i, x_i)^2 \tilde{h}(x_i, x_j)^2$ with $j \neq i$, or of the form $\tilde{h}(x_i, x_i) \tilde{h}(x_i, x_j)^3$ with $j \neq i$. In both cases there are at most $O(n)$ such terms.

In the second case, the only nonzero terms are $\tilde{h}(x_i, x_j)^4$ with $j \neq i$ (at most $O(n)$ such terms), and $\tilde{h}(x_i, x_j)^2 \tilde{h}(x_i, x_k)^2$, with unique $k, j \neq i$ (at most $O(n^2)$).

Therefore, we get

$$\mathbb{E} \left(f_*(x_i) - \tilde{f}_n(x_i) \right)^4 \lesssim \frac{n^2}{n^4} = \frac{1}{n^2},$$

which implies that

$$\begin{aligned}\mathbb{E} [C_2(\mathbf{X})^2] &\leq \sum_{i,j} \left[\mathbb{E} \left(f_*(x_i) - \tilde{f}_n(x_i) \right)^4 \right]^{1/2} \left[\mathbb{E} \left(f_*(x_j) - \tilde{f}_n(x_j) \right)^4 \right]^{1/2} \\ &\lesssim n^2 \frac{1}{n^2} = 1.\end{aligned}$$

By the Chebyshev inequality, again we have the desired bound. ■

B.2. Proofs in Section 4

Proof [Proof of Corollary 4] Here we only give a sketch of the steps to avoid repetitions. First, make approximation of the given kernel by a weighted sum of polynomial-type inner-product kernels. Note that

$$\cos \angle(\tilde{x}, \tilde{x}') = \frac{x^\top x'}{\|x\| \|x'\|},$$

which is an inner product kernel divided by $\|\tilde{x}\| \|\tilde{x}'\|$. Therefore, we write kernel k as

$$k(x, x')/d = \sum_{\iota=0}^{\infty} \alpha_\iota \left\| \frac{1}{\sqrt{d}} x \right\|^{1-\iota} \left\| \frac{1}{\sqrt{d}} x' \right\|^{1-\iota} \left(\frac{x^\top x'}{d} \right)^\iota$$

For a constant ι_0 large enough, we would have $k^{[\leq \iota_0]}$ so close to k so that

$$\left| \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 - \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} k^{[\leq \iota_0]}(\mathbf{X}, x)\|^2 \right|$$

is of the order $ne^{-\Omega(\iota_0)}$. Then we need only to upper bound

$$\mathbb{E}_x \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} k^{[\leq \iota_0]}(\mathbf{X}, x)\|^2.$$

Define kernels h_i as

$$h_i(x, x') := \left(\frac{x^\top x'}{d} \right)^i,$$

then we have

$$k(x, x') = \sum_{i=0}^{\infty} \alpha_i \left\| \frac{1}{\sqrt{d}} x \right\|^{1-\iota} \left\| \frac{1}{\sqrt{d}} x' \right\|^{1-\iota} h_i(x, x').$$

Define a diagonal matrix

$$A := \text{diag}(\|x_1/\sqrt{d}\|, \dots, \|x_n/\sqrt{d}\|).$$

Then we have

$$k^{[\iota]}(\mathbf{X}, \mathbf{X}) = \alpha_\iota A^{1-\iota} h_\iota(\mathbf{X}) A^{1-\iota}$$

and

$$k^{[\iota]}(\mathbf{X}, x) = \alpha_\iota A^{1-\iota} h_\iota(\mathbf{X}, x) \|x/\sqrt{d}\|^{1-\iota}.$$

Now we have

$$k^{[\iota]}(\mathbf{X})^{-1} k^{[\iota]}(\mathbf{X}, x) = A^{\iota-1} h_\iota(\mathbf{X}, x) \|x/\sqrt{d}\|^{1-\iota}. \quad (81)$$

Note that w.h.p., we have

$$A \asymp \mathbf{I}_n, \quad \|x/\sqrt{d}\| \asymp 1, \quad (82)$$

which implies $\|x/\sqrt{d}\|^\iota \asymp 1$ for all $\iota \leq \iota_0 = o(\sqrt{d}/\log n)$ uniformly.

Therefore as in Proposition 1, we proceed with

$$\begin{aligned} & \mathbb{E}_x \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} k^{[\leq \iota_0]}(\mathbf{X}, x)\|^2 \\ & \lesssim \sum_{i=0}^{\iota} \mathbb{E}_x \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} k^{[i]}(\mathbf{X}, x)\|^2 + \mathbb{E}_x \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} \sum_{i=\iota+1}^{\iota_0} k^{[i]}(\mathbf{X}, x)\|^2 \\ & \lesssim \sum_{i=0}^{\iota} \mathbb{E}_x \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1} k^{[i]}(\mathbf{X}, x)\|^2 + \|k^{[\leq \iota_0]}(\mathbf{X}, \mathbf{X})^{-1}\|_{\text{op}}^2 \cdot \mathbb{E}_x \left\| \sum_{i=\iota+1}^{\iota_0} k^{[i]}(\mathbf{X}, x) \right\|^2 \\ & \leq \sum_{i=0}^{\iota} \mathbb{E}_x \|k^{[i]}(\mathbf{X}, \mathbf{X})^+\|_{\text{op}}^2 \cdot \|k^{[i]}(\mathbf{X}, x)\|^2 + O(1) \frac{n}{d^{\iota+1}} \\ & \lesssim \sum_{i=0}^{\iota} \|h_i(\mathbf{X}, \mathbf{X})^+\|_{\text{op}}^2 \cdot \mathbb{E}_x \|h_i(\mathbf{X}, x)\|^2 + \frac{n}{d^{\iota+1}} \\ & \lesssim \sum_{i=0}^{\iota} (nd^{-i})^{-2} \frac{n}{d^i} + \frac{n}{d^{\iota+1}} \\ & \lesssim \frac{d^\iota}{n} + \frac{n}{d^{\iota+1}}. \end{aligned}$$

Now adding back the truncated term, for $\iota_0 = o(\sqrt{d}/\log n)$, we have

$$\mathbb{E}_{x \sim \mathcal{P}_X} \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 \lesssim \frac{d^\iota}{n} + \iota_0 \frac{n}{d^{\iota+1}} + n e^{-\Omega(\iota_0)} . \quad (83)$$

Take $\iota_0 = C(\log d + \log n) \gg \iota$ with C large enough, it is clear that $\iota_0 \ll \sqrt{d}/\log n$, we would have

$$\mathbb{E}_x \|k(\mathbf{X}, \mathbf{X})^{-1} k(\mathbf{X}, x)\|^2 \lesssim \frac{d^\iota}{n} + \iota_0 \frac{n}{d^{\iota+1}} .$$

■