
Near-Interpolators: Rapid Norm Growth and the Trade-Off between Interpolation and Generalization

Yutong Wang
MIDAS, University of Michigan

Rishi Sonthalia
UCLA

Wei Hu
University of Michigan

Abstract

We study the generalization capability of nearly-interpolating linear regressors: β 's whose training error τ is positive but small, i.e., below the noise floor. Under a random matrix theoretic assumption on the data distribution and an eigendecay assumption on the data covariance matrix Σ , we demonstrate that any near-interpolator exhibits rapid norm growth: for τ fixed, β has squared ℓ_2 -norm $\mathbb{E}[\|\beta\|_2^2] = \Omega(n^\alpha)$ where n is the number of samples and $\alpha > 1$ is the exponent of the eigendecay, i.e., $\lambda_i(\Sigma) \sim i^{-\alpha}$. This implies that existing data-independent norm-based bounds are necessarily loose. On the other hand, in the same regime we precisely characterize the asymptotic trade-off between interpolation and generalization. Our characterization reveals that larger norm scaling exponents α correspond to worse trade-offs between interpolation and generalization. We verify empirically that a similar phenomenon holds for nearly-interpolating shallow neural networks.

1 INTRODUCTION

Regularization (Nakkiran et al., 2020) and early stopping (Ji et al., 2021) are techniques to mitigate the effect of harmful overfitting by training models to *nearly*, rather than perfectly, interpolate the training data. A key question is: *how do near-interpolators generalize?*

A long line of work has investigated this question for *perfect*-interpolators. Zhang et al. (2017) noted the surprising phenomnon that, even with noise, perfect-interpolators do not necessarily overfit catastrophically,

and can still generalize to some extent. The phenomenon is later formalized as “benign overfitting” and proven to hold in linear regression in Bartlett et al. (2020). Mallinar et al. (2022) introduced the more nuanced notion of *tempered* overfitting which is closer to the empirical observation in Zhang et al. (2017) that the test error of perfect-interpolators do degrade somewhat. Koehler et al. (2021) establish a setting under which benign overfitting *can* be explained by uniform bounds. However, to the best of our knowledge, no work have studied generalization of near-interpolating linear regression ¹.

Our results. Under random matrix-theoretic and power-law spectra assumptions, we prove that nearly-interpolating ridge regressors have norms that grows rapidly (Theorem 1.6), implying that existing (non-asymptotic) generalization bounds are loose (Section 4.2). Moreover, we derive the exact formula relating the large-sample limit training and the testing error (Theorem 1.5), using the eigenlearning framework of Simon et al. (2023). Finally, we show that our theoretical results on near-interpolating ridge regressors are relevant empirically and can give insight into the behavior of early-stopped near-interpolating shallow neural networks.

Implications. Our result on the norm growth implies that existing data-independent bounds and possible extension are necessarily loose. See Section 4.2. Thus, in order to explain the learning capability of near-interpolators, there is a need to develop data-dependent generalization bounds.

On the other hand, our result allows the analysis of the trade-off between nearness of interpolation and generalization. Our result reveals delicate interplay between the overparametrization ratio and the power-law spectra exponent. In particular, for larger power-law spectra exponent implies larger asymptotic excess test error *ratio* of 5%- over 50%-noise floor interpolation, for

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹Ghosh and Belkin (2022) establishes *lower bound* for the interpolation-generalization trade-off, which is related to but distinct from our contributions.

instance. Put more simply, the harmfulness of overfitting depends on the data distribution. Moreover, this effect is stronger at high level of overparametrization (large p/n). (See Figure 2 and Figure 1-left panel.) Experimentally, this effect appears in shallow neural networks as well (Figure 4).

1.1 Related works

Near interpolation. Learning algorithms that (nearly) interpolate the training data, such as deep neural networks, have been surprisingly effective in practice despite conventional statistical wisdom suggesting otherwise (Zhang et al., 2021). Many practices in modern machine learning e.g., early stopped neural network and high-dimensional ridge regression, result in *near-* rather than *perfect-*interpolators (Ji et al., 2021; Kuzborskij and Szepesvári, 2022). In terms of theory work, Ghosh and Belkin (2022) provides a *lower bounds* on the test error for near-interpolators.

Power law spectra. Empirically, power law spectra arise in neural tangent kernels computed from practical networks for common datasets, e.g., MNIST (Veliikanov and Yarotsky, 2021) On the theory side, the power law spectra assumption has been used previously to analyze *benign* (Bartlett et al., 2020, Theorem 2) and *tempered overfitting* phenomena (Mallinar et al., 2022, Theorem 3.1).

Looseness of existing generalization bounds. Our work is motivated by the empirical evidence found by Wei et al. (2022) suggests that norms of kernel ridge regressors grow rapidly potentially beyond the purview of norm-based bound. We confirm that bounds similar to ones in Koehler et al. (2021, Corollary 1) grow to infinity. Even more refined bounds such as Koehler et al. (2021, Theorem 1) grow as n goes to infinity. Therefore, our work suggests that explaining the generalization capability of near-interpolators will require new tools.

1.2 Notations

Throughout this work, we assume the setting of high-dimensional linear regression as described below. Let n denote the number of samples and p denote the feature dimension. Consider the setting where $n, p \rightarrow \infty$ at the same time. The *sample-to-feature ratio* is denoted $\gamma := n/p \in \mathbb{R}_{>0}$ and the asymptotic sample-to-feature ratio is denoted $\gamma_* := \lim_{n \rightarrow \infty} \gamma \in \mathbb{R}_{\geq 0}$. Here, n is the fundamental parameter which p depends on implicitly.

Let $X \in \mathbb{R}^p$ and $Y \in \mathbb{R}$ denote a random vector (resp. variable), referred to as the *sample* (resp. *label*). Suppose that $\beta^* \in \mathbb{R}^p$ is such that $Y = \varepsilon + X^\top \beta^*$ where $\varepsilon \in \mathbb{R}$ is a random variable denoting independent, zero

mean noise, i.e., $\mathbb{E}[\varepsilon] = 0$ and $\varepsilon \perp X$. Here \perp denotes independence between random variables. The noise variance is denoted $\sigma^2 := \mathbb{E}[\varepsilon^2]$.

The training data is denoted $\{(x_i, y_i)\}_{i=1}^n$ where $x_i \in \mathbb{R}^p$ and $\varepsilon_i \in \mathbb{R}$ are i.i.d realizations of X and of the noise, and $y_i = \varepsilon_i + x_i^\top \beta^*$. Let $\mathbf{X} = [x_1, \dots, x_n] \in \mathbb{R}^{p \times n}$ be the data matrix obtained by horizontally stacking the x_i 's, and let $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ be the (column vector) by concatenating the y_i 's. Likewise, define $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^\top \in \mathbb{R}^n$. For a positive integer p , let \mathbf{I}_p denote the $p \times p$ identity matrix. Let $\beta \in \mathbb{R}^p$ be arbitrary. The *empirical* training mean squared error (MSE) of β is denoted

$$\mathcal{E}_{\text{train}}^n(\beta) = \frac{1}{n} \|\mathbf{X}^\top \beta - \mathbf{y}\|_2^2.$$

Likewise, the *expected* test error of β is denoted

$$\mathcal{E}_{\text{test}}^n(\beta) = \mathbb{E}[\|X^\top \beta - Y\|_2^2].$$

Let $\hat{\Sigma} := n^{-1} \mathbf{X} \mathbf{X}^\top$ denote the sample covariance matrix and $\hat{\mathbf{G}} := n^{-1} \mathbf{X}^\top \mathbf{X}$ the (scaled) gram matrix. Let $\Sigma := \mathbb{E}[\hat{\Sigma}]$ denote the population covariance.

We note that all quantities defined on the training data implicitly depend on n . When the dependencies need to be made explicit, we shall write β_n^* , $\hat{\Sigma}_n$ and so on.

1.3 Our contributions

Recall the minimum norm near-interpolator:

Definition 1.1. Let $\tau \in (0, \sigma^2)$. The *minimum norm τ -near-interpolator* is defined as

$$\beta_\tau := \operatorname{argmin}_{\beta \in \mathbb{R}^p} \|\beta\|_2^2 \text{ s.t. } \frac{1}{n} \|\mathbf{X}^\top \beta - \mathbf{y}\|_2^2 \leq \tau. \quad (1)$$

A τ -near-interpolator (not necessarily of the minimum norm) is any $\beta \in \mathbb{R}^p$ satisfying the inequality in (1).

In the overparameterized ($p > n$) regime, near-interpolators can often be realized by ridge regression:

Definition 1.2. Let $\varrho > 0$. The *ridge regressor with regularizer ϱ* is defined as

$$\hat{\beta}_\varrho := \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{n} \|\mathbf{X}^\top \beta - \mathbf{y}\|_2^2 + \varrho \|\beta\|_2^2. \quad (2)$$

Main problems: For any $\tau \in (0, \sigma^2)$, 1. find a sequence of regularizers $\{\varrho_n\}_n$ so that

$$\mathcal{E}_{\text{train}}^* := \lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{E}_{\text{train}}^n(\hat{\beta}_{\varrho_n})] = \tau \quad (3)$$

where the expectation is over all sources of randomness and 2. compute the associated asymptotic test error:

$$\mathcal{E}_{\text{test}}^* := \lim_{n \rightarrow \infty} \mathcal{E}_{\text{test}}^n(\hat{\beta}_{\varrho_n}). \quad (4)$$

Definition 1.3. Let $\tau \in (0, \sigma^2)$ be arbitrary and $\{\varrho_n\}_n$ be any sequence of regularizers. If Equation (3) holds, then we say the sequence of regressors $\{\hat{\beta}_{\varrho_n}\}_n$ is an *asymptotic τ -near interpolator* with asymptotic test error given by Equation (4).

Assumption 1.4 (Power-law spectra²). Suppose there exists $\alpha > 1$ such that the population data covariance matrix $\Sigma = \text{diag}(\lambda_1, \dots, \lambda_p)$ where $\lambda_i = i^{-\alpha}$.

There are many examples of random matrix ensembles exhibiting power-law spectra in a broader sense than that of Assumption 1.4. For instance, see (Arous and Guionnet, 2008; Mahoney and Martin, 2019; Wang et al., 2024). For simplicity, we do not pursue a general setting and will consider the setting of Assumption 1.4.

The asymptotic test error of an asymptotic τ -near interpolator can be calculated as follows. Let $F = {}_2F_1$ denote the Gaussian Hypergeometric function (Dutka, 1984, Eqn. (27)).

Theorem 1.5 (Exact trade-off formula). Let $\tau \in (0, \sigma^2)$ be arbitrary. Suppose that $\sup_{n=1,2,\dots} \|\beta^*\|_2 < +\infty$, Assumption 1.4 holds, and $X = \Sigma^{1/2}Z$ where $Z \sim \mathcal{N}(0, \mathbf{I}_p)$. There exists unique number $k \in \mathbb{R}_{>0}$ such that the following hold. Define the *regularizer-factor*

$$r := 1 - \gamma_*^{-1} F\left(1, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right)$$

and let $\varrho_n := rn^{-\alpha}$. Then $\{\hat{\beta}_{\varrho_n}\}_n$ is an asymptotic τ -near interpolator whose asymptotic test error is

$$\mathcal{E}_{\text{test}}^* = \sigma^2 \frac{1}{1 - \gamma_*^{-1} F\left(2, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right)}. \quad (5)$$

Moreover, $\mathcal{E}_{\text{test}}^*$ is a decreasing function w.r.t α , fixing all other quantities.

The reason we call Theorem 1.5 an ‘‘exact trade-off formula’’ is that Equation (5) allows the calculate of the trade-off curve between train and test error (Figure 1-right). The fundamental parameter is k . The asymptotic testing error and training error, i.e., τ , all depend on k via monotonic 1-1 correspondences on the domain $k \in (k_{\text{crit}}, \infty)$. See Figure 3 below. Thus, the asymptotic testing error depends on the training error implicitly through k .

Figure 1-right panel demonstrates that, empirically, training and test MSEs concentrate closely around Equation (5). We further discuss in detail the implications of Theorem 1.5 after stating Proposition 3.2.

Figure 1-right shows that for near-interpolators, the test error does not degrade much when training below the noise floor. A natural question is if this can

²Also referred to as the eigenvalue decay condition (Goel and Klivans, 2017).

be explained by data-independent norm-based generalization bound such as the one found in Koehler et al. (2021). Our next result shows that the growth rate of an asymptotic τ -near interpolator is superlinear:

Theorem 1.6 (Rapid norm growth). In the situation of Theorem 1.5, for any $\tau \in (0, \sigma^2)$, suppose $\varrho_n > 0$ is a sequence of regularizers such that $\{\hat{\beta}_{\varrho_n}\}_n$ is an asymptotic τ -near interpolator. Then $\mathbb{E}[\|\hat{\beta}_{\varrho_n}\|_2^2] = \Omega(n^\alpha)$.

As a consequence, data-independent norm-based generalization bound for near-interpolators, similar to the one in Koehler et al. (2021), are necessarily loose. See Section 4.2.

The key technical result that enables the proof of Theorem 1.6 is the following

Proposition 1.7 (Rapid norm growth - generic). Suppose Assumption 1.4 holds and the random matrix-theoretic Assumptions 2.5, 2.6 and 2.9 all hold. For any $r > 0$, let $\varrho_n := rn^{-\alpha}$ be the regularizer for the ridge regression. Then $\mathbb{E}[\|\hat{\beta}_{\varrho_n}\|_2^2] = \Omega(n^\alpha)$.

Remark 1.8. Proposition 1.7 still holds when the stronger Assumption 1.4 is replaced by the weaker Assumption 2.2. See Proposition 4.1.

Remark 1.9 (Effective-factor). The quantity k in Theorem 1.5 has the following interpretation. Let $\kappa := kn^{-\alpha}$, which is known as the *effective regularizer* in Wei et al. (2022). The connection between the effective regularizer and the statistical learning theoretic literature’s notion of *effective dimension* is explained in (Jacot et al., 2020a, §4.1). For this reason, we refer to k with the shortened name *eff-reg-factor*.

1.4 Organization

In Section 2, we present the necessary background as well as new technical on random matrix theory (RMT). In Section 3 and 4, we sketch the proof of Theorem 1.5 and Proposition 1.7, respectively. In Section 4.2, we discuss the implication of our results on the looseness of norm-based generalization bounds. In Section 5, we discuss our experiments. We discuss related works and the context of our work in greater details in Section 6. Finally, we conclude with discussion of future works and limitations.

2 PRIMER ON RANDOM MATRIX THEORY

We start with a fundamental concept in random matrix theory (RMT), followed by a review of RMT adapted to the power-law spectra setting and a new result (Proposition 2.10) to prepare for our main results.

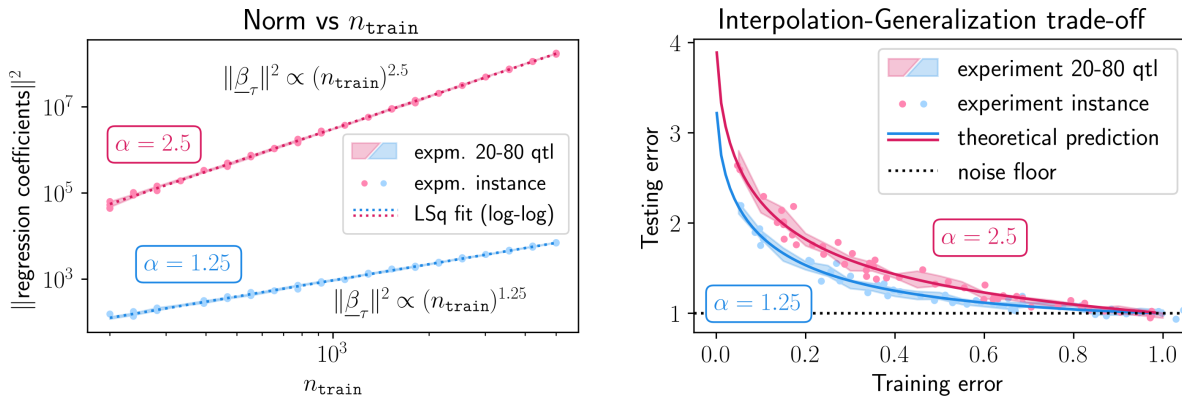


Figure 1: **Left:** *Synthetic experiments validating the norm lower bound of norms of 0.2-near-interpolators given by Proposition 1.7.* The squared norms are fitted by least squares (in log-log space) to estimate the norm-growth exponent α using only data points. See Section 5 for additional experiment details. **Right:** *Trade-off between the testing and training errors from Theorem 1.5.* The solid curves are the parametrized curves whose (x, y) -coordinates are $(\mathcal{E}_{\text{train}}^*, \mathcal{E}_{\text{test}}^*)$ and parametrized by k (which is in 1-to-1 correspondence with r see Theorem 1.5). The scatter points, subsampled for visualization, denote ridge regression run results on the HDA model (Example 2.8). The colored ribbons denote the 20-80 quantiles for the scatter points. The horizontal dotted line denotes the noise σ^2 which is set to 1 without the loss of generality.

Definition 2.1 (Empirical spectral measure). For $c \in \mathbb{R}$, let δ_c denote the Dirac measure on \mathbb{R} at c . Let $\mathbf{M} \in \mathbb{R}^{p \times p}$ be a matrix with real eigenvalues $\lambda_1, \dots, \lambda_p$. The *empirical spectral measure* of \mathbf{M} , denoted by $\text{esd}(\mathbf{M})$, is the measure on \mathbb{R} given by $\text{esd}(\mathbf{M}) = \frac{1}{p} \sum_{i=1}^p \delta_{\lambda_i}$.

Our random matrix theoretic-assumptions differs from the standard RMT ones in order to accommodate for power-law spectra. The following is a random matrix-theoretic extension of the earlier Assumption 1.4:

Assumption 2.2 (Power-law spectra, RMT version). In the situation of Section 1.2, let $\alpha > 1$ and H be some probability measure on $\mathbb{R}_{\geq 0}$. Assume that $\text{esd}(n^\alpha \Sigma)$ converges to H (in the sense of convergence in distribution). We refer to H as the α -scaled limiting spectral distribution (α -scaled LSD).

Morally, we can think of the above α as the same as that of Assumption 1.4.

Remark 2.3 (Comparison with standard LSD). In RMT, the condition that “ $\text{esd}(\Sigma)$ converges to H ” is standard, where H is simply referred to as the *limiting spectral distribution* (LSD) (Bai and Silverstein, 2010). For power-law spectra covariance, i.e., Σ satisfying Assumption 1.4, $\text{esd}(\Sigma)$ may not have a measure-theoretic limit while $\text{esd}(n^\alpha \Sigma)$ does, as we will show in Section 4.1.

Definition 2.4 (Stieltjes transform). Let μ be a measure on \mathbb{R} with support S . The *Stieltjes transform* of μ is the (complex-valued) function with input $z \in \mathbb{C} \setminus S$ given by $\mathcal{S}_\mu(z) := \int \frac{\mu(t) dt}{t-z}$.

Next, we recall the so-called the *self-consistent equa-*

tion (Tao, 2011) which relates the regularizer-factor r with the eff-reg-factor k :

Assumption 2.5 (Self-consistent equation). In the situation of Section 1.2, denote by λ_i the i -th largest eigenvalue of Σ . For each $r > 0$, there exists a unique $k \equiv k(r) \in \mathbb{R}$ such that the tuple (r, k) satisfies

$$1 = \frac{r}{k} + \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^p \frac{1}{1 + kn^{-\alpha} \lambda_i^{-1}}. \quad (6)$$

Similar to Remark 2.3, Equation (6) includes a $n^{-\alpha}$ scaling term that does not show up in the standard self-consistent equation. Again, our Assumption 2.5 differs from this standard one in order to deal with the power-law spectra.

Next, we state a version of the classical Marchenko-Pastur law for a random matrix \mathbf{X} (and its associated Gram matrix $n^\alpha \check{\mathbf{G}}_n$):

Assumption 2.6 (Marchenko-Pastur law). In the setting of Assumption 2.5, further assume that

$$\lim_{n \rightarrow \infty} r \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}}_n)}(-r) = k \mathcal{S}_H(-k), \quad \text{almost surely}$$

and $\lim_{n \rightarrow \infty} \frac{d}{dr} \left(r \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}}_n)}(-r) \right) = \frac{d}{dr} (k \mathcal{S}_H(-k))$. We note that the k on the RHS depends on r .

Remark 2.7. Assumption 2.5 and Assumption 2.6 are standard assumptions in random matrix theory. Both of them are satisfied by the well-studied *high-dimensional asymptotic* (HDA) model (Example 2.8). For instance, see Dobriban and Wager (2018) under “Marchenko-Pastur theorem”.

The HDA model serves as an exemplary model in random matrix theory possessing many properties that are particularly amenable to analysis. It is defined as:

Example 2.8. Let $\gamma_* \in (0, \infty)$. The *high-dimensional asymptotic (HDA)*³ model:

1. $\mathbf{X} = \Sigma^{1/2}\mathbf{Z}$ where the entries of $\mathbf{Z} = \{Z_{ij}\} \in \mathbb{R}^{p \times n}$ are i.i.d, have zero mean $\mathbb{E}[Z_{ij}] = 0$ and unit variance $\mathbb{E}[Z_{ij}^2] = 1$. The matrix Σ is positive semidefinite.
2. $n/p \rightarrow \gamma_*$
3. Spectral distribution of $n^\alpha \Sigma$ converges to a distribution H supported on $\mathbb{R}_{\geq 0}$.

Note that our Example 2.8 is somewhat different compared to the *conventional* HDA model, wherein the third item is “ Σ converges to a distribution H ”. Since we are working with power-law spectra in the covariance matrix, we require the n^α coefficient in our Example 2.8.

We now state a new random matrix-theoretic assumption that is one of the key steps for proving rapid norm growth under the RMT setting (Proposition 4.1):

Assumption 2.9 (Positivity condition). In the setting of Assumption 2.5, further assume that for every $r > 0$, we have

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{d}{dr} (r \mathcal{S}_{\text{esd}(n^\alpha \mathbf{G}_n)}(-r)) \right] > 0.$$

We show that the HDA model satisfies Assumption 2.9, a fact that appears to be new:

Proposition 2.10. Assumption 2.9 holds for the HDA model.

We prove the proposition in Appendix C. Now, having introduced the necessary RMT background, we now turn to proving Theorem 1.5 on the interpolation-generalization trade-off.

3 INTERPOLATION-GENERALIZATION TRADE-OFF

Simon et al. (2023) derived “estimates” of the testing and training errors of kernel ridge regression. These estimates, dubbed the *eigenlearning framework*, are non-rigorous⁴ due to invoking a *Gaussian universality* condition. However, when the kernel is linear and

³See Bai and Silverstein (2010); Dobriban and Wager (2018).

⁴See Mallinar et al. (2022) for a thorough discussion. Works in similar vein include Bordelon et al. (2020); Canatar et al. (2021)

the data is Gaussian (as is the case in Theorem 1.5), the framework *is* rigorous. See Jacot et al. (2020b).

Given this, we use the eigenlearning framework to rigorously calculate the asymptotic training and testing error of the estimators in Theorem 1.5. To this end, we first define two key functions of the eff-reg-factor k (See Remark 1.9 for the terminology):

Definition 3.1. Let $\alpha > 1$ and $\gamma_* \in [0, \infty)$. Define functions⁵ $\mathcal{I}(\cdot) \equiv \mathcal{I}_{\alpha, \gamma_*}(\cdot)$ and $\mathcal{J}(\cdot) \equiv \mathcal{J}_{\alpha, \gamma_*}(\cdot)$ as

$$\mathcal{I}(k) := \int_0^{\frac{1}{\gamma_*}} \frac{dx}{1 + kx^\alpha}, \quad \mathcal{J}(k) := \int_0^{\frac{1}{\gamma_*}} \frac{dx}{(1 + kx^\alpha)^2}.$$

When $\gamma_* = 0$, we take $1/0 := +\infty$.

These integrals arise in explicit calculations of the eigenlearning equations for the train and test MSEs applied to our setting. They can be computed via the integral representation of the Gaussian hypergeometric function given in (Dutka, 1984, Eqn. (27)). The calculations are in Appendix A.1, where we show that \mathcal{I} and \mathcal{J} are, respectively, equal to

$$\int_0^{\frac{1}{\gamma_*}} \frac{dx}{1 + kx^\alpha} = \gamma_*^{-1} F\left(1, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right), \text{ and}$$

$$\int_0^{\frac{1}{\gamma_*}} \frac{dx}{(1 + kx^\alpha)^2} = \gamma_*^{-1} F\left(2, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right).$$

To relate the above to $\mathcal{E}_{\text{test}}^*, \mathcal{E}_{\text{train}}^*$ estimates of the testing and training errors in the eigenlearning framework, we prove

Proposition 3.2. In the situation of Theorem 1.5,

$$\mathcal{E}_{\text{test}}^* \equiv \lim_{n \rightarrow \infty} \mathcal{E}_{\text{test}}^n(\hat{\beta}_\ell) = \sigma^2 \cdot \frac{1}{1 - \mathcal{J}(k)}, \quad \text{and}$$

$$\mathcal{E}_{\text{train}}^* \equiv \lim_{n \rightarrow \infty} \mathbb{E}[\mathcal{E}_{\text{train}}^n(\hat{\beta}_\ell)] = \sigma^2 \cdot \frac{(1 - \mathcal{I}(k))^2}{1 - \mathcal{J}(k)}.$$

Moreover, there exists $k_{\text{crit}} \in \mathbb{R}_{\geq 0}$ such that

1. For each $r > 0$, there exists a unique $k \in (k_{\text{crit}}, +\infty)$ such that $r = \mathcal{R}(k) := k(1 - \mathcal{I}(k))$,
2. \mathcal{R} is monotonically increasing on $(k_{\text{crit}}, +\infty)$,
3. $\mathcal{E}_{\text{test}}^* > \sigma^2$ for all $k \in (k_{\text{crit}}, +\infty)$,
4. $\lim_{k \rightarrow +\infty} \mathcal{E}_{\text{test}}^* = \sigma^2$.

For the proof of Proposition 3.2, see Appendix A.2. At a high level, to prove the first part we apply the eigenlearning framework while accounting for the additional layer of complexity due to the power-law spectra. For

⁵Throughout this work, α and γ_* are fixed constants. For brevity, we often simply write \mathcal{I} or \mathcal{J} .




Figure 2: *Larger power-law spectra exponent implies larger asymptotic excess test error when interpolating to 5% of the noise compared to 50% of the noise floor.* Let $\mathcal{E}_{\text{test}}^* = \mathcal{E}_{\text{test}}^*(\alpha, \gamma_*, \tau)$ be as in Theorem 1.5 where we make the dependency on parameters α, γ_*, τ explicit. The color and contour line of plot shows the ratio of test errors at two levels of nearness of interpolation $\mathcal{E}_{\text{test}}^*(\alpha, \gamma_*, 0.05)/\mathcal{E}_{\text{test}}^*(\alpha, \gamma_*, 0.5)$ over an (α, γ_*) -grid.

the ‘‘Moreover’’ part, we directly analyze $\mathcal{R}, \mathcal{E}_{\text{train}}^*$ and $\mathcal{E}_{\text{test}}^*$ as functions of r, k and α . Now, note that Proposition 3.2 immediately implies Theorem 1.5.

We now discuss some of the consequences of Proposition 3.2 First, \mathcal{R} is a bijection that relates the eff-reg-factor k and the regularizer-factor r . The plot of \mathcal{R} is visualized in Figure 3. Furthermore, note that $\lim_{k \rightarrow +\infty} \mathcal{E}_{\text{test}}^* = \sigma^2$ precisely states that the test error can be made arbitrarily close to the noise floor as k (equivalently, r) goes to infinity.

Using Proposition 3.2 with the implementation of ${}_2F_1$ in SCIPY, we illustrate the trade-off between the training error versus the test error in Figure 1-Right and the test error ratio in Figure 2.

Remark 3.3 (Data-independent regularizer-selection). Let $\tau \in (0, \sigma^2)$ be a desired level of nearness of interpolation. To select a regularizer ϱ_n that achieves τ' -near-interpolation for $\tau' \approx \tau$, we use the following method: *Step 1.* Find $k \in (k_{\text{crit}}, +\infty)$ such that $\mathcal{E}_{\text{train}}^* = \mathcal{E}_{\text{train}}^*(k) = \tau$, using the expression for $\mathcal{E}_{\text{train}}^*(k)$ in Proposition 3.2. Let k_τ be such a k . *Step 2.* Next, set $r := \mathcal{R}(k_\tau)$, where \mathcal{R} is also as in Proposition 3.2. *Step 3.* Set the regularizer as $\varrho_n := rn^\alpha$.

4 RAPID NORM GROWTH

Proposition 1.7 can be proven in even greater generality under random matrix-theoretic assumptions.

Proposition 4.1 (Rapid norm growth under RMT). Suppose Assumptions 2.2, 2.5, 2.6, and 2.9 all hold. For any $r > 0$, let $\varrho_n := rn^{-\alpha}$ be the regularizer for

the ridge regression. Then $\mathbb{E}[\|\hat{\beta}_{\varrho_n}\|_2^2] = \Omega(n^\alpha)$.

The goal of this section is to sketch the proof for Proposition 4.1. Complete proofs of all results are included in the Appendix. Throughout, we assume the setting of Section 1.2.

The next step is the following:

Proposition 4.2. Let $\varrho := rn^{-\alpha}$. Then we have

$$\mathbb{E}\|\hat{\beta}_\varrho\|_2^2 \geq n^\alpha \sigma^2 \cdot \mathbb{E}\left[\frac{d}{dr}(r\mathcal{S}_{\text{esd}(n^\alpha \mathbf{\check{G}})}(-r))\right].$$

The proof of Proposition 4.2 and other omitted proofs in this section can be found in Appendix B.

Given Proposition 4.2, the proof of Proposition 4.1 is straightforward:

Proof of Proposition 4.1. Let

$$L := \lim_{n \rightarrow \infty} \mathbb{E}\left[\frac{d}{dr}(r\mathcal{S}_{\text{esd}(n^\alpha \mathbf{\check{G}})}(-r))\right] > 0$$

be as in Assumption 2.9. Thus, for all $n \gg 0$ sufficiently large, we have $\mathbb{E}\left[\frac{d}{dr}(r\mathcal{S}_{\text{esd}(n^\alpha \mathbf{\check{G}})}(-r))\right] > L/2 > 0$. By Proposition 4.2, we get that $\mathbb{E}\|\hat{\beta}_\varrho\|_2^2 \geq n^\alpha \sigma^2 \cdot \frac{L}{2}$ for all $n \gg 0$, as desired. \square

4.1 α -scaled limiting spectral distribution

In this section, we check that power-law spectra covariance matrices has an α -scaled limiting spectral distribution. In other words, Assumption 1.4 implies Assumption 2.2. Note that this is necessary because we have proved Proposition 4.1 rather than Proposition 1.7. So we need to make sure that the special case, i.e., Proposition 1.7, is indeed the ‘‘special case’’.

Definition 4.3. Given a measure μ on \mathbb{R} , we let $\text{cdf}[\mu]$ denote the cumulative distribution function (CDF) of μ .

Proposition 4.4. Under Assumption 1.4, we have that Assumption 2.2 holds. In other words,

$$\lim_{n \rightarrow \infty} \text{cdf}[\text{esd}(n^\alpha \Sigma)](t) = \begin{cases} 1 - \gamma_* t^{-1/\alpha} & : t \geq \gamma_*^\alpha \\ 0 & \text{otherwise.} \end{cases}$$

4.2 Looseness of norm-based generalization bounds

Conjecturally, a norm-based generalization bound for near-interpolators should have the following form: under suitable assumptions, with high probability

$$\sup_{\beta: \|\beta\| \leq B, \mathcal{E}_{\text{train}}(\beta) \leq \tau} \mathcal{E}_{\text{test}}(\beta) = O\left(\frac{B^2 \text{Tr}(\Sigma)}{n}\right).$$




Figure 3: The $\mathcal{R}(k)$ function from Proposition 3.2. The x -axis is the input k . Note that for $k < k_{\text{crit}}$ the regularizer r is negative. Although we are only interested in the $(k_{\text{crit}}, +\infty)$ portion, negative regularizers have been studied by Tsigler and Bartlett (2020); Wu and Xu (2020).

where $\tau \in \mathbb{R}_{\geq 0}$. For $\tau = 0$, the best known bound for (perfect) interpolators is given by Koehler et al. (2021, Corollary 1) under Gaussian assumption on the data $X \sim \mathcal{N}(0, \Sigma)$ and $B \geq \|\beta^*\|$.

To the best of our knowledge, there is no known extension to the case of near-interpolators, i.e., where $\tau > 0$. However, such bound is not informative for our scenario, since by Theorem 1.5 and Proposition 1.7, it is possible to choose ϱ_n such that $\underline{1}$. $\mathbb{E}[\mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho_n})] \rightarrow \tau$, $\underline{2}$. $\mathcal{E}_{\text{test}}(\hat{\beta}_{\varrho_n}) \rightarrow c \in \mathbb{R}_{\geq 0}$, and $\underline{3}$. $\|\beta_{\varrho_n}\|^2 = \Omega(n^\alpha)$ for any $\alpha > 0$. Thus, the bound goes to infinity while $\mathcal{E}_{\text{test}}(\hat{\beta}_{\varrho_n})$ is finite.

5 EXPERIMENTS

We run two types of synthetic experiments. The first type, plotted in Figure 1, employs (linear) ridge regression. The second type, plotted in Figure 4 employs neural networks. The data for both types of experiments are drawn from the HDA model (Example 2.8). Moreover, we have conducted experiments on several real world regression datasets from the UCI regression collection.

5.1 Experiments on synthetic data

To generate Figure 1-left, we run experiments with $\alpha \in \{1.25, 2.5\}$ and $\frac{n}{p} = \gamma_* = \frac{2}{3}$. We sweep over the train MSE parameter $\mathcal{E}_{\text{train}}^*$ to explore the trade-off between the train and testing MSE in linear ridge regression as described in Theorem 1.5. The value of $\mathcal{E}_{\text{train}}^*$ are swepted on a linearly-spaced grid of size 16 from $0.05\sigma^2$ to $0.8\sigma^2$. The parameters are $n_{\text{train}} =$

5000, $n_{\text{test}} = 1000$, $\gamma_* = 0.5$, $\alpha = 1.75$ and $\sigma^2 = 1$.

The regularizer achieving a desired training MSE $\tau \in (0, \sigma^2)$ is chosen according to the method described in Remark 3.3. We sample $\beta^* \in \mathbb{R}^p$ such that β_i^* are i.i.d Gaussian with zero mean and variance $= 10/p$.

The same set up is used for Figure 1-right, except we sweep over n_{train} rather than over $\mathcal{E}_{\text{train}}^*$. The value of n_{train} are swepted on a logarithmically-spaced grid of size 20 from 200 to 5000.

For Figure 4, the identical set-up is used as in the Figure 1, except ridge regression is replaced with neural networks⁶ during training. We emphasize that the ground truth data (i.e., the teacher) is still generated via the same linear function β^* . All code for the experiments are included in Appendix E.

Remark 5.1. As discussed in the introduction, near-interpolating neural network and its interpolation-generalization trade-offs exhibit similar phenomenon as in the case of ridge regression. Namely, larger power-law spectra exponent implies larger asymptotic excess test error when interpolating to 5% of the noise compared to 50% of the noise floor, for instance.

Remark 5.2. Intriguingly, in both the ridge regression (Figure 1) and neural network (Figure 4) experiments, the setting with to larger value of the norm-growth exponent α (right panels) results in poorer trade-offs (left panels). For ridge regression, this is explained by the “Moreover” part of Theorem 1.5. We believe it is an interesting future direction whether this behavior can be proved in the neural network setting.

5.2 Experiments on UCI datasets

We conduct experiments on the **forest** and the **stock** datasets from the UCI regression collection (Kelly et al., 2023). Using neural tangent kernels (Arora et al., 2019), we observe power-law spectra in the kernel matrices for both of these datasets (Figure 5-right and Figure 6-right). In this subsection, we discuss the experiments on the **forest** dataset in relation to our theoretical results. Due to space constraints, we refer the reader to Appendix F for details of the experimental setup.

In Figure 5-left, note that the curve corresponding to **forest.2-1** has the fastest spectra decay and simultaneously the worst trade-off. Evidently, larger decay exponent corresponds to a poorer trade-off, especially for near-interpolators, i.e., as the training error approaches 0. This is in agreement with our theoretical results under random matrix theory assumptions illustrated in Figure 2. Similar phenomenon occurs for the

⁶We use the default settings in sklearn, except with early stopping for near-interpolation.




Figure 4: Experiments with neural networks (top row: 1-hidden layer, bottom row: 5-hidden layers). Analogous to Figure 1. See Section 5 and Remark 5.1 for details.

stock dataset. See Appendix F.

6 ADDITIONAL RELATED WORKS AND NOVELTY OF OUR WORK

Technical Novelty of Theorem 1.7. Prior works of Hastie et al. (2022); Dereziński et al. (2019) require a lower bound on the smallest eigenvalue of the covariance matrix. Hence, the scenario studied in this paper is not amenable to their results. Further, Theorem 1.5 shows that when we have a sharper decay (larger α) of the eigenvalues, we have a worse tradeoff between the test and training error. Hence, we show that the scenario from Hastie et al. (2022) is the most benign one. This well conditioning assumption is relaxed in Cheng and Montanari (2022), however, they require that $\|\Sigma^{-1/2}\beta\|$ is finite. We do not require this assumption. Finally, Dobriban and Wager (2018) does not need the well-conditioning assumption but instead needs to assume an isotropic distribution on β . Since we work with fixed β , their results do not apply.

Trade-offs in interpolation-based learning. Prior works (Ghosh and Belkin, 2022; Belkin et al., 2018; Sonthalia et al., 2023) have studied the tradeoff between near interpolation and generalization. For regression, pre-

vious works have also studied the fundamental trade-off in learning algorithms between overparametrization and (Lipschitz) smoothness (Bubeck and Sellke, 2021), and robustness and smoothness (Zhang et al., 2022).

Loosens of existing generalization bounds. Belkin et al. (2018, Theorem 1) establishes for *classification* that the RKHS norm of a “near-interpolating” classifier grows at rate $\Omega(\exp(n^{1/p}))$. The growth is unbounded if $n = \Omega(\exp(p))$. If the number of samples $n = \Theta(\text{poly}(p))$, then the lower bound does not grow to infinity. While our results are for regression and thus not directly comparable, our lower bound is meaningful in the more practical $n \propto p$ regime.

Power-law spectra datasets. Synthetic data with artificial power law EVD covariance have been used frequently as toy examples (Berthier et al., 2020; Mallinar et al., 2022). On real datasets, power law EVD is often observed to describe neural tangent kernels (NTK) well in practice, including on MNIST ((Bahri et al., 2021, Fig. 4) and (Velikanov and Yarotsky, 2022, Fig. 2)), FASHION-MNIST (Cui et al., 2021, Fig. 7) CALTECH 101 (Murray et al., 2022, Fig. 1), CIFAR-100 (Wei et al., 2022, Fig. 3).

Theoretical machine learning works using power-law spectra. Bordelon et al. (2020) shows that power law




Figure 5: *Left.* Training/testing error trade-off on the “forest” dataset from the UCI regression dataset collection using kernel ridge regression with the neural tangent kernel. Each curve is labeled by “DatasetName.d-f” where “d” and “f” represents the number of layers and the number of *fixed* layers in the NTK corresponding to ReLU networks. *Right.* The eigenvalue index vs eigenvalue plot of the NTK matrix exhibits power-law spectra. A tiny value is added to the eigenvalues for better visualization on the log-scale.

EVD implies power law learning curve. Velikanov and Yarotsky (2021, §6.2) computes the power law EVD exponent for certain NTKs with ReLU to be $\alpha = 1 + \frac{1}{d}$. Murray et al. (2022) computes the EVD for NTKs with several different activations. Bartlett et al. (2020, Theorem 6) shows that benign overfitting occurs when the covariance matrix eigenvalues $\lambda_i = i^{-1} \log^{-b}(i+1)$ for $b > 1$. Mallinar et al. (2022) studies power law decay for $\alpha \geq 1$ and proposes a taxonomy of overfitting into three categories: catastrophic, tempered and benign. The EVD condition is also known as the *capacity condition* in the kernel ridge regression literature. See Bietti et al. (2021) and the references there-in.

Random matrix theory (RMT). The signal processing research community have long been using RMT for theoretical analysis (Couillet and Debbah, 2012). Increasingly RMT has been applied to machine learning as well as a key tool for analysis.

In particular, Dobriban and Wager (2018); Hastie et al. (2022); Jacot et al. (2020b); Liang and Rakhlin (2020) have applied RMT for (kernel) ridge regression, Sonthalia and Nadakuditi (2023); Kausik et al. (2023) use it to understand generalization of linear denoisers, Paquette et al. (2022, 2021) uses the so-called local Marchenko-Pastur law (Knowles and Yin, 2017) to analyze gradient-based algorithms. Finally, Wei et al. (2022) also applies such local law to analyze the so-called *generalized cross-validation (GCV) estimator*.

7 DISCUSSION AND LIMITATIONS

We conclude with several future research directions that we believe will be fruitful:

Connection to early stopping. Typically, early stopping prevents the trained algorithm from perfectly

interpolating the data. Can early stopped learning theory results, e.g., Ji et al. (2021); Kuzborskij and Szepesvári (2022), be applied to analyze near-interpolators?

Near-interpolators and uniform convergence generalization bound. Is possible to use uniform convergence-based approach to give non-vacuous generalization bound under the setting studied in this work? This question has already been raised by Dobriban and Wager (2018) in the context of classification in a similar setting. An interesting question is if classical learning theory can be used to obtain results that are currently only obtained via random matrix theoretic or similar techniques. Another approach is to extend the results of Koehler et al. (2021) to the near-interpolation setting.

Limitations. Our work is restricted to analyzing a random matrix model. Understanding the phenomenon uncovered in this paper in more general models and additional real world settings will be needed. Moreover, our work does not rule out the existence of uniform convergence generalization bound.

Code availability. Code used to run and plot the experiments shown in all figures is available at <https://github.com/YutongWangUMich/Near-Interpolators-Figures>.

Acknowledgements

YW acknowledges support from the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Futures program. WH acknowledges support from the Google Research Scholar program.

References

- Arora, S., Du, S. S., Li, Z., Salakhutdinov, R., Wang, R., and Yu, D. (2019). Harnessing the power of infinitely wide deep nets on small-data tasks. In *International Conference on Learning Representations*.
- Arous, G. B. and Guionnet, A. (2008). The spectrum of heavy tailed random matrices. *Communications in Mathematical Physics*, 278(3):715–751.
- Bahri, Y., Dyer, E., Kaplan, J., Lee, J., and Sharma, U. (2021). Explaining neural scaling laws. *arXiv preprint arXiv:2102.06701*.
- Bai, Z. and Silverstein, J. W. (2010). *Spectral analysis of large dimensional random matrices*, volume 20. Springer.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070.
- Belkin, M., Ma, S., and Mandal, S. (2018). To understand deep learning we need to understand kernel learning. In *International Conference on Machine Learning*, pages 541–549. PMLR.
- Berthier, R., Bach, F., and Gaillard, P. (2020). Tight nonparametric convergence rates for stochastic gradient descent under the noiseless linear model. *Advances in Neural Information Processing Systems*, 33:2576–2586.
- Bietti, A., Venturi, L., and Bruna, J. (2021). On the sample complexity of learning with geometric stability. In *Advances in Neural Information Processing Systems*.
- Bordelon, B., Canatar, A., and Pehlevan, C. (2020). Spectrum dependent learning curves in kernel regression and wide neural networks. In *International Conference on Machine Learning*, pages 1024–1034. PMLR.
- Bubeck, S. and Sellke, M. (2021). A universal law of robustness via isoperimetry. *Advances in Neural Information Processing Systems*, 34:28811–28822.
- Canatar, A., Bordelon, B., and Pehlevan, C. (2021). Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks. *Nature communications*, 12(1):1–12.
- Cheng, C. and Montanari, A. (2022). Dimension free ridge regression. *arXiv preprint arXiv:2210.08571*.
- Couillet, R. and Debbah, M. (2012). Signal processing in large systems: A new paradigm. *IEEE Signal Processing Magazine*, 30(1):24–39.
- Cui, H., Loureiro, B., Krzakala, F., and Zdeborová, L. (2021). Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. In *Advances in Neural Information Processing Systems*, pages 10131–10143.
- Derezinski, M., Liang, F. T., and Mahoney, M. W. (2019). Exact expressions for double descent and implicit regularization via surrogate random design. *ArXiv*, abs/1912.04533.
- Dobriban, E. and Wager, S. (2018). High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279.
- Dutka, J. (1984). The early history of the hypergeometric function. *Archive for History of Exact Sciences*, pages 15–34.
- Ghosh, N. and Belkin, M. (2022). A universal trade-off between the model size, test loss, and training loss of linear predictors. *arXiv preprint arXiv:2207.11621*.
- Goel, S. and Klivans, A. (2017). Eigenvalue decay implies polynomial-time learnability for neural networks. *Advances in Neural Information Processing Systems*, 30.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022). Surprises in High-Dimensional Ridgeless Least Squares Interpolation. *The Annals of Statistics*.
- Jacot, A., Simsek, B., Spadaro, F., Hongler, C., and Gabriel, F. (2020a). Implicit regularization of random feature models. In *International Conference on Machine Learning*, pages 4631–4640. PMLR.
- Jacot, A., Simsek, B., Spadaro, F., Hongler, C., and Gabriel, F. (2020b). Kernel alignment risk estimator: Risk prediction from training data. *Advances in Neural Information Processing Systems*, 33:15568–15578.
- Ji, Z., Li, J., and Telgarsky, M. (2021). Early-stopped neural networks are consistent. *Advances in Neural Information Processing Systems*, 34:1805–1817.
- Karp, D. B. and López, J. L. (2017). Representations of hypergeometric functions for arbitrary parameter values and their use. *Journal of Approximation Theory*, 218:42–70.
- Kausik, C., Srivastava, K., and Sonthalia, R. (2023). Generalization error without independence: Denoising, linear regression, and transfer learning. *arXiv preprint arXiv:2305.17297*.
- Kelly, M., Longjohn, R., and Nottingham, K. (2023). The uci machine learning repository. <https://archive.ics.uci.edu>.
- Knowles, A. and Yin, J. (2017). Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169(1):257–352.

- Koehler, F., Zhou, L., Sutherland, D. J., and Srebro, N. (2021). Uniform convergence of interpolators: Gaussian width, norm bounds and benign overfitting. *Advances in Neural Information Processing Systems*, 34:20657–20668.
- Kuzborskij, I. and Szepesvári, C. (2022). Learning lipschitz functions by gd-trained shallow overparameterized relu neural networks. *arXiv preprint arXiv:2212.13848*.
- Liang, T. and Rakhlin, A. (2020). Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 48(3):1329–1347.
- Mahoney, M. and Martin, C. (2019). Traditional and heavy tailed self regularization in neural network models. In *International Conference on Machine Learning*, pages 4284–4293. PMLR.
- Mallinar, N. R., Simon, J. B., Abedsoltan, A., Pandit, P., Belkin, M., and Nakkiran, P. (2022). Benign, tempered, or catastrophic: Toward a refined taxonomy of overfitting. In *Advances in Neural Information Processing Systems*.
- Murray, M., Jin, H., Bowman, B., and Montufar, G. (2022). Characterizing the spectrum of the NTK via a power series expansion. *arXiv preprint arXiv:2211.07844*.
- Nakkiran, P., Venkat, P., Kakade, S. M., and Ma, T. (2020). Optimal regularization can mitigate double descent. In *International Conference on Learning Representations*.
- Paquette, C., Lee, K., Pedregosa, F., and Paquette, E. (2021). Sgd in the large: Average-case analysis, asymptotics, and stepsize criticality. In *Conference on Learning Theory*, pages 3548–3626. PMLR.
- Paquette, C., van Merriënboer, B., Paquette, E., and Pedregosa, F. (2022). Halting time is predictable for large models: A universality property and average-case analysis. *Foundations of Computational Mathematics*, pages 1–77.
- Silverstein, J. W. and Choi, S.-I. (1995). Analysis of the limiting spectral distribution of large dimensional random matrices. *Journal of Multivariate Analysis*, 54(2):295–309.
- Simon, J. B., Dickens, M., Karkada, D., and Dewese, M. (2023). The eigenlearning framework: A conservation law perspective on kernel ridge regression and wide neural networks. *Transactions on Machine Learning Research*.
- Sonthalia, R., Li, X., and Gu, B. (2023). Underparameterized double descent for ridge regularized least squares denoising of data on a line. *arXiv preprint arXiv:2305.14689*.
- Sonthalia, R. and Nadakuditi, R. R. (2023). Training data size induced double descent for denoising feedforward neural networks and the role of training noise. *Transactions on Machine Learning Research*.
- Tao, T. (2011). Intuitive understanding of the Stieltjes transform. MathOverflow. Version: 2011-10-25.
- Tsigler, A. and Bartlett, P. L. (2020). Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*.
- Velikanov, M. and Yarotsky, D. (2021). Explicit loss asymptotics in the gradient descent training of neural networks. *Advances in Neural Information Processing Systems*, 34:2570–2582.
- Velikanov, M. and Yarotsky, D. (2022). Tight convergence rate bounds for optimization under power law spectral conditions. *arXiv preprint arXiv:2202.00992*.
- Wang, Z., Engel, A., Sarwate, A. D., Dumitriu, I., and Chiang, T. (2024). Spectral evolution and invariance in linear-width neural networks. *Advances in Neural Information Processing Systems*, 36.
- Wei, A., Hu, W., and Steinhardt, J. (2022). More than a toy: Random matrix models predict how real-world neural representations generalize. In *Proceedings of the 39th International Conference on Machine Learning*, pages 23549–23588. PMLR.
- Wu, D. and Xu, J. (2020). On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems*, 33:10112–10123.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2021). Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.
- Zhang, H., Wu, Y., and Huang, H. (2022). How many data are needed for robust learning? *arXiv preprint arXiv:2202.11592*.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Not Applicable]

- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
- (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix

A Proof of Theorem 1.5 — the exact trade-off formula

Our goal is to calculate the asymptotic test error $\mathcal{E}_{\text{test}}^*$ under the assumptions of Theorem 1.5. This is accomplished through the following three steps.

The first step is to calculate the closed-form solution for the integrals defined in Definition 3.1 which are key ingredients for the expression of $\mathcal{E}_{\text{test}}^*$. This is done in Appendix A.1. The second step is to relate the integrals from Definition 3.1 to the self-consistent equations in Equation (6). This is done in Appendix A.2. The final step is to relate the self-consistent equations Equation (6) to the asymptotic test error $\mathcal{E}_{\text{test}}^*$. This is done in Appendix A.3.

A.1 Closed-form expression for the integrals in Definition 3.1

We prove the identities

$$\int_0^{\frac{1}{\gamma_*}} \frac{dx}{1+kx^\alpha} = \gamma_*^{-1} F\left(1, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right), \text{ and}$$

$$\int_0^{\frac{1}{\gamma_*}} \frac{dx}{(1+kx^\alpha)^2} = \gamma_*^{-1} F\left(2, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right).$$

as shown in the main text following Definition 3.1.

Proof. Let ${}_2F_1(a, b; c; z)$ be the *Gauss hypergeometric function*. Note that the function can be implemented in SCIPY as `scipy.special.hyp2f1` and is used to plot Figure 1. Let Γ denote the Gamma function. The integral representation of the Gaussian hypergeometric function is well-known and is given by⁷

$${}_2F_1(\sigma, a; b; -z) = \frac{\Gamma(b)}{\Gamma(a)\Gamma(b-a)} \int_0^1 \frac{t^{a-1}(1-t)^{b-a-1}}{(1+zt)^\sigma} dt. \quad (7)$$

Moreover, (7) is finite for $z \in \mathbb{C} \setminus (-\infty, -1]$ and $\Re(b-a) > 0$ and $\Re(a) > 0$. Thus, by Equation (7), we have

$${}_2F_1\left(1, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right) = \frac{\Gamma(1 + \frac{1}{\alpha})}{\Gamma(\frac{1}{\alpha})\Gamma(1)} \int_0^1 \frac{t^{1/\alpha} t^{-1}}{1+k\gamma_*^{-\alpha}t} dt = \frac{1}{\alpha} \int_0^1 \frac{t^{1/\alpha} t^{-1}}{1+k\gamma_*^{-\alpha}t} dt$$

where the second inequality follows from the well-known identity $z = \Gamma(1+z)/\Gamma(z)$ for the Gamma function. Let $u = t^{1/\alpha}$. Then we have $du = \frac{1}{\alpha} t^{1/\alpha} t^{-1} dt$. Thus, by u-substitution, we have

$$\frac{1}{\alpha} \int_0^1 \frac{t^{1/\alpha} t^{-1}}{1+k\gamma_*^{-\alpha}t} dt = \int_0^1 \frac{1}{1+k\gamma_*^{-\alpha}u^\alpha} du = \int_0^{1/\gamma_*} \frac{\gamma_*}{1+kx^\alpha} dx$$

where the second inequality used u-substituted with $x = u/\gamma_*$. Now, we have by the definition of \mathcal{I} in Definition 3.1 that

$$\mathcal{I}(k) = \int_0^{1/\gamma_*} \frac{1}{1+kx^\alpha} dx = \frac{1}{\gamma_*} \int_0^{1/\gamma_*} \frac{\gamma_*}{1+kx^\alpha} dx = \frac{1}{\gamma_*} {}_2F_1\left(1, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right)$$

as desired. By an analogous calculation, we get $\mathcal{J}(k) = \frac{1}{\gamma_*} {}_2F_1\left(2, \frac{1}{\alpha}; 1 + \frac{1}{\alpha}; -k\gamma_*^{-\alpha}\right)$. □

A.2 Proof of Proposition 3.2

We begin by analyzing the functions defined in Definition 3.1 and prove the items 1 and 2 of the “Moreover” part of Proposition 3.2:

⁷We used the formula stated in Karp and López (2017).

Proposition A.1. Let \mathcal{I} and \mathcal{J} be functions as defined in Definition 3.1. Under Assumption 1.4 and Assumption 2.5, we have that $r = \mathcal{R}(k) := k \cdot (1 - \mathcal{I}(k))$ and $\frac{dr}{dk} = 1 - \mathcal{J}(k)$.

Furthermore, the following holds:

1. $\mathcal{R}(k) \asymp k$ for $k \gg 0$,
2. There exists $k_{\text{crit}} > 0$ such that $\mathcal{R}(k_{\text{crit}}) = 0$, \mathcal{R} is increasing and positive on $(k_{\text{crit}}, +\infty)$.
3. $\mathcal{J}(k) < 1$ for $k \in (k_{\text{crit}}, +\infty)$ and $\mathcal{J}(+\infty) = 0$.

Proof of Proposition A.1. We begin by proving the first part: that $r = \mathcal{R}(k) := k \cdot (1 - \mathcal{I}(k))$ and $\frac{dr}{dk} = 1 - \mathcal{J}(k)$. Rewrite the limit in Equation (6) as follows:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^p \frac{1}{1 + kn^{-\alpha} \lambda_i^{-1}} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n/\gamma} \frac{1}{1 + k(i/n)^\alpha} = \int_0^{1/\gamma_*} \frac{dx}{1 + kx^\alpha}$$

The right-most equality follows from the definition of the (Riemann) integral. If $\gamma_* = 0$, then $1/\gamma_* = +\infty$ and the above is interpreted as an improper Riemann integral. Now, rearranging Equation (6), we get the desired formula of $r = \mathcal{R}(k) := k \cdot (1 - \mathcal{I}(k))$. The formula for $\frac{dr}{dk}$ follows by “differentiating under the integral” (Leibniz integral rule).

For the first item of the “Furthermore” part, it suffices to show that $\lim_{k \rightarrow +\infty} \mathcal{I}(k) = 0$. This follows from the fact that $\lim_{k \rightarrow +\infty} \frac{1}{1 + kx^\alpha} = 0$ for all $x > 0$, integrability of the function $(1 + x^\alpha)^{-1}$ over $\mathbb{R}_{\geq 0}$, and the dominated convergence theorem. Likewise, $\lim_{k \rightarrow \infty} \mathcal{J}(k) = 0$ as well.

For the second item of the “Furthermore” part, we note that for all x sufficiently large, we have $\frac{dr}{dk} > 0$ since $\lim_{k \rightarrow \infty} \mathcal{J}(k) = 0$. Now, let k_{crit} be the largest real number such that $\mathcal{R}(k_{\text{crit}}) = 0$. Since $\mathcal{R}(0) = 0$, we must have $k_{\text{crit}} \geq 0$.

For all $k > k_{\text{crit}}$, we claim that $\mathcal{I}(k) < 1$. To see this, assume the contrary. Then by the fact that $\lim_{k \rightarrow +\infty} \mathcal{I}(k) = 0$ and the intermediate value theorem, there must exist $k' > k$ such that $\mathcal{I}(k') = 1$ which implies that $\mathcal{R}(k') = 0$. This contradicts the maximality of k_{crit} .

Finally, since $1 + kx^\alpha \leq (1 + kx^\alpha)^2$ for all $k \geq 0$ and $x \geq 0$, we have that $\mathcal{I}(k) \geq \mathcal{J}(k)$ for all such k 's. Thus, by the previous claim, for all $k > k_{\text{crit}}$, we have $1 > \mathcal{I}(k) \geq \mathcal{J}(k)$. This proves that $\frac{dr}{dk} > 0$ for all $k > k_{\text{crit}}$, as desired. \square

A.3 Review of the eigenlearning framework (Simon et al., 2023)

Before proceeding with finishing the proof of Proposition 3.2, we briefly review the eigenlearning framework. Simon et al. (2023) calculates the test error for the estimator

$$\tilde{\beta}_\delta := \mathbf{X}(\mathbf{X}^\top \mathbf{X} + \delta \mathbf{I}_n)^{-1} \mathbf{y} = \mathbf{X}(n\tilde{\mathbf{G}} + \delta \mathbf{I}_n)^{-1} \mathbf{y} \quad (8)$$

for kernel ridge regression using the so-called *eigenlearning equations* (Simon et al., 2023, Section 4.1). Below, we recall some relevant parts of the framework:

Definition A.2 (Eigenlearning eqn. specialized to setting in Section 1.3). Suppose that the ground truth regression function is linear, i.e., $f(x) = x^\top \beta^*$ for some $\beta^* \in \mathbb{R}^p$. Let δ and κ satisfy the equation

$$n = \frac{\delta}{\kappa} + \sum_{i=1}^p \frac{\lambda_i}{\lambda_i + \kappa}. \quad (9)$$

Define the following n -dependent quantities:

1. *Overfitting coefficient*: $\mathcal{E}_{\text{coef}} := n \frac{d\kappa}{d\delta}$
2. *Testing error*: $\mathcal{E}_{\text{test}} := \mathcal{E}_{\text{coef}}(\sigma^2 + C)$ where

$$C = \sum_{i=1}^p (1 - \mathcal{L}_i)(\beta_i^*)^2 \quad \text{and} \quad \mathcal{L}_i := \frac{\lambda_i}{\lambda_i + \kappa}.$$

3. *Training error:* $\mathcal{E}_{\text{train}} := \frac{\delta^2}{n^2 \kappa^2} \mathcal{E}_{\text{test}}$.

Proof of Proposition 3.2. Simon et al. (2023) uses a different scaling for ridge regression than the one we use. To bridge the different notations, we first resolve this discrepancy. Comparing Equation (8) with the expression in Equation (11), if we let $\delta := n\varrho$, then the expressions are equivalent, i.e., $\check{\beta}_\delta = \hat{\beta}_\varrho$. To see this, note that

$$\begin{aligned} \check{\beta}_\delta &= \check{\beta}_{n\varrho} = X(X^\top X + n\varrho \mathbb{I}_n)^{-1} y \\ &= (X X^\top + n\varrho \mathbb{I}_p)^{-1} X y \quad \because \text{Lemma B.2} \\ &= (n(n^{-1} X X^\top + \varrho \mathbb{I}_p))^{-1} X y \\ &= (\hat{\Sigma} + \varrho \mathbb{I}_p)^{-1} \frac{1}{n} X y = \hat{\beta}_\varrho \quad \because \text{Definition of } \hat{\beta}_\varrho \end{aligned}$$

Below, let $r > 0$ be arbitrary. Furthermore, we claim that as $n \rightarrow \infty$, we have r, k satisfies Equation (6) if and only if the tuple $(\delta, \kappa) := (nrn^{-\alpha}, kn^{-\alpha})$ satisfies Equation (9):

$$n = \frac{\delta}{\kappa} + \sum_{i=1}^p \frac{\lambda_i}{\lambda_i + \kappa} \iff n = \frac{nrn^{-\alpha}}{kn^{-\alpha}} + \sum_{i=1}^p \frac{\lambda_i}{\lambda_i + kn^{-\alpha}} \iff 1 = \frac{r}{k} + \frac{1}{n} \sum_{i=1}^p \frac{1}{1 + kn^{-\alpha} \lambda_i^{-1}}.$$

Taking limit as $n \rightarrow \infty$, we have proved the claim.

Next, we show that $\lim_{n \rightarrow \infty} C = 0$ where C is as in Definition A.2. We have $\mathcal{L}_i := \frac{\lambda_i}{\lambda_i + \kappa} = \frac{1}{1 + k(i/n)^\alpha}$. Note that $\lim_{n \rightarrow \infty} \mathcal{L}_i = 1$ for all fixed i . On the other hand, since $\sup_{n=1,2,\dots} \|\beta^*\|_2 < +\infty$, dominated convergence theorem implies that $\lim_{n \rightarrow \infty} C = 0$.

We claim that the following asymptotic expression for the testing and training error hold:

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{test}} = \sigma^2 \cdot \frac{dk}{dr} \quad \text{and} \quad \lim_{n \rightarrow \infty} \mathcal{E}_{\text{train}} = \sigma^2 \cdot \frac{r^2}{k^2} \cdot \frac{dk}{dr} \quad (10)$$

where r and k satisfy Equation (6) from Assumption 2.5.

To see this, first note that the overfitting coefficient satisfies

$$\mathcal{E}_{\text{coef}} := n \frac{d\kappa}{d\delta} = n \frac{d\kappa}{d\varrho} \frac{d\varrho}{d\delta} = n \frac{d\kappa}{d\varrho} \frac{1}{n} = \frac{d\kappa}{d\varrho} = \frac{dk}{dr}.$$

Thus, we obtain the following asymptotic expression

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{test}} = \mathcal{E}_{\text{coef}} \cdot \sigma^2 = \sigma^2 \cdot \frac{dk}{dr}.$$

On the other hand, the training error is given by

$$\mathcal{E}_{\text{train}} = \frac{\delta^2}{n^2 \kappa^2} \mathcal{E}_{\text{test}} = \frac{\varrho^2}{\kappa^2} \mathcal{E}_{\text{test}} = \mathcal{E}_{\text{test}} \cdot \frac{r^2}{k^2}.$$

Therefore, $\lim_{n \rightarrow \infty} \mathcal{E}_{\text{train}} = \sigma^2 \cdot \frac{r^2}{k^2} \cdot \frac{dk}{dr}$. This proves (10), as desired. \square

B Proof of Proposition 4.1 — rapid norm growth under RMT assumptions

The first key technical step the following:

Proposition B.1. $\mathbb{E} \|\hat{\beta}_\varrho\|_2^2 \geq n^{-1} \sigma^2 \mathbb{E}[\text{tr}((\hat{\Sigma} + \varrho \mathbf{I}_p)^{-2} \hat{\Sigma})]$.

Proof. Below, for brevity we let $\mathbf{a} := \mathbf{X}^\top \beta^*$ and $\mathbf{M} := (\hat{\Sigma} + \varrho \mathbf{I}_p)^{-1} \frac{1}{n} \mathbf{X}$. Recall the closed-form solution for Equation (2) is given by the formula

$$\hat{\beta}_\varrho := (\hat{\Sigma} + \varrho \mathbf{I}_p)^{-1} \frac{1}{n} \mathbf{X} y. \quad (11)$$

Thus,

$$\hat{\beta}_\varrho = (\hat{\Sigma} + \varrho \mathbf{I}_p)^{-1} \frac{1}{n} \mathbf{X} y = (\hat{\Sigma} + \varrho \mathbf{I}_p)^{-1} \frac{1}{n} \mathbf{X} (f(\mathbf{X}) + \varepsilon) = \mathbf{M}(\mathbf{a} + \varepsilon).$$

Thus,

$$\|\hat{\boldsymbol{\beta}}_\varrho\|_2^2 = (\mathbf{a} + \boldsymbol{\varepsilon})^\top \mathbf{M}^\top \mathbf{M} (\mathbf{a} + \boldsymbol{\varepsilon}) = \underbrace{\mathbf{a}^\top \mathbf{M}^\top \mathbf{M} \mathbf{a}}_{\geq 0} + \boldsymbol{\varepsilon}^\top \mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon} + 2\boldsymbol{\varepsilon}^\top \mathbf{M}^\top \mathbf{M} \mathbf{a}.$$

Note that $\boldsymbol{\varepsilon} \perp \mathbf{M}^\top \mathbf{M} \mathbf{a}$ since $\boldsymbol{\varepsilon} \perp \mathbf{X}$. Thus, since $\mathbb{E}[\boldsymbol{\varepsilon}] = 0$, we have

$$\mathbb{E}[\|\hat{\boldsymbol{\beta}}_\varrho\|_2^2] = \mathbb{E}[(\mathbf{a} + \boldsymbol{\varepsilon})^\top \mathbf{M}^\top \mathbf{M} (\mathbf{a} + \boldsymbol{\varepsilon})] \geq \mathbb{E}[\boldsymbol{\varepsilon}^\top \mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon}] = \mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top)].$$

Since $\mathbf{M}^\top \mathbf{M} \perp \boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top$, we have

$$\mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top)] = \text{tr}(\mathbb{E}[\mathbf{M}^\top \mathbf{M}] \mathbb{E}[\boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top]) = \text{tr}(\mathbb{E}[\mathbf{M}^\top \mathbf{M} \sigma^2 \mathbf{I}_n]) = \sigma^2 \mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M})].$$

On the other hand, $\mathbf{M}^\top \mathbf{M} = \frac{1}{n}(\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p)^{-1} \hat{\boldsymbol{\Sigma}} (\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p)^{-1}$. Using the cyclic property of trace, we get the desired inequality. \square

Proof sketch of Prop. B.1. We first simplify $\|\hat{\boldsymbol{\beta}}_\varrho\|_2^2$ using the well-known formula for ridge regression: Next, let $\mathbf{M} := (\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p)^{-1} \frac{1}{n} \mathbf{X}$. Using the independence of \mathbf{X} and $\boldsymbol{\varepsilon}$, we get $\mathbb{E}[\|\hat{\boldsymbol{\beta}}_\varrho\|_2^2] \geq \mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top)]$. Since $\mathbf{M}^\top \mathbf{M}$ and $\boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top$ are also independent, we have

$$\mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M} \boldsymbol{\varepsilon} \boldsymbol{\varepsilon}^\top)] = \sigma^2 \mathbb{E}[\text{tr}(\mathbf{M}^\top \mathbf{M})].$$

By $\mathbf{M}^\top \mathbf{M} = \frac{1}{n}(\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p)^{-1} \hat{\boldsymbol{\Sigma}} (\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p)^{-1}$ and the cyclic property of trace, we get the desired inequality. \square

For the sake of completeness, we prove Equation (11) though its well-known

Proof of Equation (11). Start with the objective function $\mathcal{F}(\boldsymbol{\beta}) := \frac{1}{n} \|\mathbf{X}^\top \boldsymbol{\beta} - \mathbf{y}\|_2^2 + \varrho \|\boldsymbol{\beta}\|_2^2$. Take derivative with respect to $\boldsymbol{\beta}$, we have

$$\begin{aligned} \frac{1}{2} \nabla_{\boldsymbol{\beta}} \left(\frac{1}{n} \|\mathbf{X}^\top \boldsymbol{\beta} - \mathbf{y}\|_2^2 + \varrho \|\boldsymbol{\beta}\|_2^2 \right) &= \frac{1}{2} \nabla_{\boldsymbol{\beta}} \left(\boldsymbol{\beta}^\top (\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p) \boldsymbol{\beta} - \frac{2}{n} \boldsymbol{\beta}^\top \mathbf{X} \mathbf{y} \right) \\ &= (\hat{\boldsymbol{\Sigma}} + \varrho \mathbf{I}_p) \boldsymbol{\beta} - \frac{1}{n} \mathbf{X} \mathbf{y}. \end{aligned}$$

Since $\nabla_{\boldsymbol{\beta}} \mathcal{F}(\hat{\boldsymbol{\beta}}_\varrho) = 0$, we are done. \square

Lemma B.2 (Special case of Woodbury formula). Let

$$\mathbf{M} \in \mathbb{R}^{p \times n} \tag{12}$$

be an arbitrary matrix and $\varrho \in (0, \infty)$. Then

$$(\mathbf{M} \mathbf{M}^\top + \varrho \mathbf{I}_p)^{-1} \mathbf{M} = \mathbf{M} (\mathbf{M}^\top \mathbf{M} + \varrho \mathbf{I}_n)^{-1} \in \mathbb{R}^{n \times p}.$$

Proof of Lemma B.2. It suffices to prove Lemma B.2 for the special case when $\varrho = 1$, which we assume below. By the Woodbury matrix identity, we have

$$(\mathbf{M} \mathbf{M}^\top + \mathbf{I}_p)^{-1} = \mathbf{I} - \mathbf{M} (\mathbf{M}^\top \mathbf{M} + \mathbf{I}_n)^{-1} \mathbf{M}^\top \tag{13}$$

For brevity, let $\mathbf{P} := \mathbf{M} \mathbf{M}^\top + \mathbf{I}_p$ and let $\mathbf{N} := \mathbf{M}^\top \mathbf{M} + \mathbf{I}_n$. To proceed, we have

$$\begin{aligned} &\mathbf{P}^{-1} \mathbf{M} \\ &= \mathbf{M} - \mathbf{M} \mathbf{N}^{-1} \mathbf{M}^\top \mathbf{M} \quad \because \text{Multiplying (13) by } \mathbf{M} \text{ on the right} \\ &= \mathbf{M} (\mathbf{I}_n - \mathbf{N}^{-1} \mathbf{M}^\top \mathbf{M}) \quad \because \text{Factoring out } \mathbf{M} \text{ on the left} \\ &= \mathbf{M} (\mathbf{I}_n - (\mathbf{I}_n - \mathbf{N}^{-1})) \quad \because \mathbf{I}_n = \mathbf{N}^{-1} \mathbf{N} = \mathbf{N}^{-1} + \mathbf{N}^{-1} \mathbf{M}^\top \mathbf{M} \\ &= \mathbf{M} \mathbf{N}^{-1} \end{aligned}$$

as desired. \square

Lemma B.3. Let $\mathbf{M} \in \mathbb{R}^{p \times p}$ be any symmetric matrix and $z \in \mathbb{R}$. Then we have

$$\frac{d}{dz} \operatorname{tr}(z(\mathbf{M} + z\mathbf{I}_p)^{-1}) = \operatorname{tr}(\mathbf{M}(\mathbf{M} + z\mathbf{I}_p)^{-2}).$$

Proof of Lemma B.3. Without the loss of generality, suppose that $\mathbf{M} = \operatorname{diag}(\lambda_1, \dots, \lambda_p)$. Then we have $f(z) := \operatorname{tr}(z(\mathbf{M} + z\mathbf{I}_p)^{-1}) = \sum_{i=1}^p \frac{z}{\lambda_i + z}$. Now, from elementary calculus, we have

$$\frac{d}{dx} \frac{x}{y+x} = (y+x)^{-1} - x(y+x)^{-2} = (y+x)^{-2}((y+x) - x) = \frac{y}{(y+x)^2}.$$

From this, we recover the fact that $\frac{d}{dz} f(z) = \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + z)^2} = \operatorname{tr}(\mathbf{M}(\mathbf{M} + z\mathbf{I}_p)^{-2})$, as desired. \square

Lemma B.4 (Gram-to-covariance). Let $c \in \mathbb{R}$ and $z \in \mathbb{C}$ be arbitrary, then $\mathcal{S}_{\operatorname{esd}(c\hat{\Sigma})}(z) = \gamma \cdot \mathcal{S}_{\operatorname{esd}(c\check{\mathbf{G}})}(z) - \frac{(1-\gamma)}{z}$.

Proof of Lemma B.4. Without the loss of generality, we may assume that $c = 1$. Let $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_p$ be the eigenvalues of $\hat{\Sigma}$. Since $p > n$, we necessarily have that $\hat{\lambda}_{n+1} = \dots = \hat{\lambda}_p = 0$. Moreover, $\hat{\lambda}_1, \dots, \hat{\lambda}_n$ are the eigenvalues of $\check{\mathbf{G}}$. Now, unwinding the definition, we have

$$\mathcal{S}_{\operatorname{esd}(\hat{\Sigma})}(z) = \frac{1}{p} \sum_{i=1}^p \frac{1}{\hat{\lambda}_i - z}$$

and

$$\mathcal{S}_{\operatorname{esd}(\check{\mathbf{G}})}(z) = \frac{1}{n} \sum_{i=1}^n \frac{1}{\hat{\lambda}_i - z}.$$

Thus,

$$\begin{aligned} \mathcal{S}_{\operatorname{esd}(\hat{\Sigma})}(z) &= \frac{1}{p} \left(\sum_{i=1}^n \frac{1}{\hat{\lambda}_i - z} + \sum_{i=n+1}^p \frac{1}{-z} \right) \\ &= \left(\frac{n}{p} \frac{1}{n} \sum_{i=1}^n \frac{1}{\hat{\lambda}_i - z} \right) - \frac{p-n}{p} \frac{1}{z} \\ &= \gamma \cdot \mathcal{S}_{\operatorname{esd}(\check{\mathbf{G}})}(z) - \frac{(1-\gamma)}{z} \end{aligned}$$

as desired. \square

Proof of Proposition 4.2. Recall from Proposition B.1 that $\mathbb{E}\|\hat{\beta}\|_2^2 \geq n^{-1}\sigma^2\mathbb{E}[\operatorname{tr}((\hat{\Sigma} + \varrho\mathbf{I}_p)^{-2}\hat{\Sigma})]$. Below, we analyze the term inside the expectation. By the definition of the Stieltjes transform, we have

$$\operatorname{tr}(\varrho(\hat{\Sigma} + \varrho\mathbf{I}_p)^{-1}) = \operatorname{tr}(rn^{-\alpha}(\hat{\Sigma} + rn^{-\alpha}\mathbf{I}_p)^{-1}) = \operatorname{tr}(r(n^\alpha\hat{\Sigma} + r\mathbf{I}_p)^{-1}) = pr\mathcal{S}_{\operatorname{esd}(n^\alpha\hat{\Sigma})}(-r).$$

Therefore, by Lemma B.3, we have

$$\begin{aligned} \frac{d}{dr} \left(pr\mathcal{S}_{\operatorname{esd}(n^\alpha\hat{\Sigma})}(-r) \right) &= \frac{d}{dr} \operatorname{tr}(\varrho(\hat{\Sigma} + \varrho\mathbf{I}_p)^{-1}) \\ &= \frac{d\varrho}{dr} \cdot \frac{d}{d\varrho} \operatorname{tr}(\varrho(\hat{\Sigma} + \varrho\mathbf{I}_p)^{-1}) = n^{-\alpha} \operatorname{tr}((\hat{\Sigma} + \varrho\mathbf{I}_p)^{-2}\hat{\Sigma}). \end{aligned}$$

By Lemma B.4, we have

$$\begin{aligned} pr\mathcal{S}_{\operatorname{esd}(n^\alpha\hat{\Sigma})}(-r) &= pr \left(\gamma \cdot \mathcal{S}_{\operatorname{esd}(n^\alpha\check{\mathbf{G}})}(-r) + \frac{(1-\gamma)}{r} \right) \\ &= nr\mathcal{S}_{\operatorname{esd}(n^\alpha\check{\mathbf{G}})}(-r) + p(1-\gamma) \end{aligned}$$

Thus, we have

$$\frac{d}{dr} \left(pr \mathcal{S}_{\text{esd}(n^\alpha \hat{\Sigma})}(-r) \right) = n \frac{d}{dr} \left(r \mathcal{S}_{\text{esd}(n^\alpha \hat{\mathbf{G}})}(-r) \right)$$

from which we conclude that

$$\text{tr}((\hat{\Sigma} + \varrho \mathbf{I}_p)^{-2} \hat{\Sigma}) = n^{\alpha+1} \frac{d}{dr} \left(r \mathcal{S}_{\text{esd}(n^\alpha \hat{\mathbf{G}})}(-r) \right).$$

In view of $\mathbb{E} \|\hat{\beta}\|_2^2 \geq n^{-1} \sigma^2 \mathbb{E}[\text{tr}((\hat{\Sigma} + \varrho \mathbf{I}_p)^{-2} \hat{\Sigma})]$ from Proposition B.1, we get the desired inequality. \square

B.1 Proof of Proposition 4.4

Proof of Proposition 4.4. To simplify notations in this proof, we write γ instead of γ_* . Now, the set of eigenvalues of $n^\alpha \Sigma$ can be expressed as

$$\begin{aligned} & \{(n/i)^\alpha\}_{i=1, \dots, p} \\ &= \underbrace{\left\{ \left(\frac{n}{p}\right)^\alpha, \dots, \left(\frac{n}{n+1}\right)^\alpha \right\}}_{=\gamma^\alpha}, \underbrace{\left\{ \frac{n}{n}, \left(\frac{n}{n-1}\right)^\alpha, \dots, \left(\frac{n}{1}\right)^\alpha \right\}}_{=n^\alpha}. \end{aligned}$$

Thus, $\text{cdf}[\text{esd}(n^\alpha \Sigma)](t) = 0$ if $t < \gamma^\alpha$ and $= 1$ if $t \geq n^\alpha$. It remains to calculate $\text{cdf}[\text{esd}(n^\alpha \Sigma)](t) = 0$ if $t < \gamma^\alpha$ for $t \in [\gamma^\alpha, n^\alpha]$.

To this end, let $t \in [\gamma^\alpha, n^\alpha)$ and $j(t) \in \{1, \dots, p\}$ be the smallest index such that $(n/j(t))^\alpha \leq t$. By definition of the CDF, we have $\text{cdf}[\text{esd}(n^\alpha \Sigma)](t) = \frac{1}{p}(p - j(t) + 1)$. We first argue that $j(t) \neq 1$ by contradiction. If $j(t) = 1$, then we have $n^\alpha \leq t$. Since $t \in [\gamma^\alpha, n^\alpha]$, this implies that $t = n^\alpha$, a contradiction. Thus, $j(t) \neq 1$.

Now, by the definition of $j(t)$, we have $(n/(j(t) - 1))^\alpha > t$. Therefore, $n/j(t) \leq t^{1/\alpha} < n/(j(t) - 1)$ which implies that $j(t) - 1 < nt^{-1/\alpha} \leq j(t)$. By the definition of the ceiling function, we have that $j(t) = \text{ceil}(nt^{-1/\alpha})$. Therefore,

$$\begin{aligned} & \text{cdf}[\text{esd}(n^\alpha \Sigma)](t) \\ &= \frac{1}{p}(p - \text{ceil}(nt^{-1/\alpha}) + 1) \\ &= 1 - \frac{\gamma}{n} \text{ceil}(nt^{-1/\alpha}) + \frac{\gamma}{n}. \end{aligned}$$

Taking limit of both side as $n \rightarrow \infty$ and using the fact that $\lim_{n \rightarrow \infty} \text{ceil}(nc)/n = c$ for any positive number $c > 0$, we get the desired result. \square

C Proof of the Positivity Condition portion of Proposition 2.10

This section will focus on the proof of Proposition 2.10, in particular the Positivity Condition (Assumption 2.9) portion. Thus, throughout this section, we assume the setting of Example 2.8. As mentioned in the main text, it is well-known that Assumptions 2.5 to 2.6 for the HDA model. As such, we assume these assumptions. Now, using Assumption 2.6 and elementary calculus, we first show that

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{d}{dr} (r \mathcal{S}_{\text{esd}(n^\alpha \hat{\mathbf{G}})}(-r)) \right] = \left(\frac{dr}{dk} \right)^{-1} \cdot \frac{d}{dk} (k \mathcal{S}_H(-k))$$

where r and k are as in Assumption 2.5. Thus, we reduce to showing the positivity of $\frac{dr}{dk}$ and $\frac{d}{dk} (k \mathcal{S}_H(-k))$.

Before proceeding, we recall several definitions and notations adapted from Dobriban and Wager (2018):

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\mathcal{S}_{\text{esd}(n^\alpha \hat{\mathbf{G}})}(z) \right] = v(z) \tag{14}$$

is analogous to the $v(z)$ defined in the paragraph immediately following (Dobriban and Wager, 2018, Eqn. (2)). The difference is our Equation (14) is for the limit of the n^α -scaled matrices $n^\alpha \hat{\mathbf{G}}$, rather than for $\hat{\mathbf{G}}$ as in Dobriban and Wager (2018).

Let $H = \lim_{n \rightarrow \infty} \text{cdf}[\text{esd}(n^\alpha \Sigma)]$ be the limiting distribution as in Assumption 2.2. Plugging in $z = -r$ into Dobriban and Wager (2018, Eqn. (A.1)), we have

$$-\frac{1}{v(-r)} = -r - \frac{1}{\gamma} \int \frac{tdH(t)}{1 + tv(-r)}.$$

Letting $k \equiv k(r) := \frac{1}{v(-r)}$, we can rewrite the above as

$$1 = \frac{r}{k} + \frac{1}{\gamma} \int \frac{tdH(t)}{k + t}. \quad (15)$$

By construction, we have

$$\frac{1}{\gamma} \int \frac{tdH(t)}{k + t} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^p \frac{1}{1 + kn^{-\alpha} \lambda_i^{-1}}$$

where the RHS is as in Assumption 2.5. Consequently, the tuple r, k from Assumption 2.5 coincide with the earlier definition of $k := \frac{1}{v(-r)}$ right before Equation (15). Having established the above, we now proceed to:

Lemma C.1. Under the HDA model (Example 2.8) and the EVD condition (Assumption 2.2), we have that $\lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{d}{dr} (r \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r)) \right] > 0$.

Proof of Lemma C.1. By the product rule, we have

$$\frac{d}{dr} \left(r \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r) \right) = \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}})}(r) - r \mathcal{S}'_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r)$$

Now, taking the limit of the above equation on both side, we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{d}{dr} \left(r \mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r) \right) \right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E} \left[\mathcal{S}_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r) - r \mathcal{S}'_{\text{esd}(n^\alpha \check{\mathbf{G}})}(-r) \right] \\ &= v(-r) - r v'(-r) \quad \because \text{Definition of } v \text{ and } v' \\ &= \frac{d}{dr} (rv(-r)) \quad \because \text{Product rule} \\ &= \frac{d}{dr} (k \mathcal{S}_H(-k)) \quad \because \text{Marchenko-Pastur law (Assumption 2.6)} \\ &= \frac{dk}{dr} \cdot \frac{d}{dk} (k \mathcal{S}_H(-k)) \quad \because \text{Chain rule} \\ &= \left(\frac{dr}{dk} \right)^{-1} \cdot \frac{d}{dk} (k \mathcal{S}_H(-k)) \quad \because \text{Inverse function theorem} \end{aligned}$$

To complete the proof, it suffices to show that both $\frac{dr}{dk}$ and $\frac{d}{dk} (k \mathcal{S}_H(-k))$ are positive which will be checked in the next two lemmas. \square

Lemma C.2. The function $\frac{dr}{dk}$ evaluated at k is positive.

Proof of Lemma C.2. Recall that $k = \frac{1}{v(-r)}$. Thus, we have

$$\frac{dk}{dr}(r) = (-1) \frac{1}{v(-r)^2} (-1) \cdot v'(-r) = \frac{v'(-r)}{v(-r)^2}.$$

From the proof of Silverstein and Choi (1995, Theorem 4.1), we see that $v'(\cdot) > 0$ for all negative inputs. In particular, $v'(-r) > 0$ which implies that $\frac{dk}{dr}$ is positive. By the inverse function theorem, we have $\frac{dr}{dk} = \left(\frac{dk}{dr} \right)^{-1}$ is also positive. \square

Lemma C.3. The quantity $\frac{d}{dk} (k \mathcal{S}_H(-k))$ is positive.

Proof of Lemma C.3. Plugging in $z = -r$ into Dobriban and Wager (2018, Eqn. (3)), we have

$$v(-r) - \frac{1}{r} = \frac{1}{\gamma} \left(m(-r) - \frac{1}{r} \right). \quad (16)$$

Now,

$$rm(-r) = \gamma rv(-r) + (1 - \gamma) \quad \because \text{Equation (16)} \tag{17}$$

$$= \gamma \frac{r}{k} + (1 - \gamma) \quad \because \text{Definition of } k \tag{18}$$

$$= \left(\gamma - \int \frac{tdH(t)}{k+t} \right) + (1 - \gamma) \quad \because \text{Equation (15)} \tag{19}$$

$$= 1 - \int \frac{tdH(t)}{k+t} \tag{20}$$

$$= \int \frac{k dH(t)}{k+t} \quad \because 1 = \int dH(t) = \int \frac{k+t}{k+t} dH(t) \tag{21}$$

$$= k \mathcal{S}_H(-k). \tag{22}$$

Thus, differentiating under the integral, we have

$$\frac{d}{dk}(k \mathcal{S}_H(-k)) = \int \frac{d}{dk} \left(\frac{k}{k+t} \right) dH(t) = \int \frac{tdH(t)}{(k+t)^2} > 0$$

as desired. □

D Proof of Theorem 1.6 — rapid norm growth

Proof of Theorem 1.6. Let $\tau' := \frac{\tau + \sigma^2}{2}$. Then by definition, we have $\tau' \in (0, \sigma^2)$. From Theorem 1.5, we can pick $r' > 0$ so that the sequence of regularizers $\varrho'_n = r' n^{-\alpha}$ satisfies $\mathcal{E}_{\text{train}}^*(r') = \tau'$. Next, we note that there exists $c > 0$ such that

$$\Pr(\mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho'_n}) \geq \mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho_n})) \geq c$$

for all n sufficiently large. Such a c is guaranteed to exist by the fact that

$$\lim_n \mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho'_n}) = \tau' = \frac{\tau + \sigma^2}{2} > \tau = \lim_n \mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho_n}).$$

Now, note that

$$\mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho'_n}) \geq \mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho_n}) \tag{23}$$

$$\iff \hat{\beta}_{\varrho_n} \text{ is a } \mathcal{E}_{\text{train}}(\hat{\beta}_{\varrho'_n})\text{-near interpolator} \tag{24}$$

$$\implies \|\hat{\beta}_{\varrho_n}\|^2 \geq \|\hat{\beta}_{\varrho'_n}\|^2 \tag{25}$$

From this, we conclude that there exists a $c > 0$ such that

$$\Pr(\|\hat{\beta}_{\varrho_n}\|^2 \geq \|\hat{\beta}_{\varrho'_n}\|^2) \geq c$$

for all n sufficiently large. From this, it follows that $\mathbb{E}[\|\hat{\beta}_{\varrho_n}\|^2] \geq c \mathbb{E}[\|\hat{\beta}_{\varrho'_n}\|^2]$. By Proposition 1.7 and the preceding inequality, we are done. □

E Code for implementation \mathcal{I} and \mathcal{J}

Implementation of the \mathcal{I} and \mathcal{J} functions from Definition 3.1 can be implemented in SCIPY as:

```

1 import scipy.special as sc
2 gamma = 0.5
3 alpha = 1.75
4
5 # I helper
6 I_gen = lambda x,k, alpha: x*sc.hyp2f1(1,(1/alpha), 1 + (1/alpha), -k*x**alpha)
7 # J helper
8 J_gen = lambda x,k, alpha: x*sc.hyp2f1(2,(1/alpha), 1 + (1/alpha), -k*x**alpha)
9

```

```

10 I = lambda k : I_gen(1/gamma, k, alpha) #\mathcal{I}
11 J = lambda k : J_gen(1/gamma, k, alpha) #\mathcal{J}
12
13 N = lambda k : 1 - I(k) # helper
14 D = lambda k : 1 - J(k) # helper
15
16 Etst = lambda k : 1/D(k) #\Etest/\sigma^2
17 Etrn = lambda k : N(k)**2/D(k) #\Etrain/\sigma^2
18 R = lambda k : k*(1-I(k)) # \mathcal{R}

```

F Experiments

Code for reproducing all figures are included in the official GitHub repository:

<https://github.com/YutongWangUMich/Near-Interpolators-Figures/>

For downloading the UCI regression datasets, we use the following repository:

https://github.com/treforevans/uci_datasets

For the neural tangent kernel, we use the official repository associated to Arora et al. (2019):

<https://github.com/LeoYu/neural-tangent-kernel-UCI>

In Figure 6-left, note that the curve corresponding to `stock.2-1` has the fastest spectra decay and simultaneously the worst trade-off. Evidently, larger decay exponent corresponds to a poorer trade-off, especially for near-interpolators, i.e., as the training error approaches 0. This is in agreement with our theoretical results under random matrix theory assumptions illustrated in Figure 2.




Figure 6: *Left*. Training/testing error trade-off on the “stock” dataset from the UCI regression dataset collection using kernel ridge regression with the neural tangent kernel. Each curve is labeled by “DatasetName.d-f” where “d” and “f” represents the number of layers and the number of *fixed* layers in the NTK corresponding to ReLU networks. *Right*. The eigenvalue index vs eigenvalue plot of the NTK matrix exhibits power-law spectra. A tiny value is added to the eigenvalues for better visualization on the log-scale.