

A Universal Trade-off Between the Model Size, Test Loss, and Training Loss of Linear Predictors

Nikhil Ghosh
 Department of Statistics
 University of California, Berkeley
 nikhil_ghosh@berkeley.edu

Mikhail Belkin
 Halıcıoğlu Data Science Institute
 University of California, San Diego
 mbelkin@ucsd.edu

Abstract

In this work we establish an algorithm and distribution independent non-asymptotic trade-off between the model size, excess test loss, and training loss of linear predictors. Specifically, we show that models that perform well on the test data (have low excess loss) are either “classical” – have training loss close to the noise level, or are “modern” – have a much larger number of parameters compared to the minimum needed to fit the training data exactly.

We also provide a more precise asymptotic analysis when the limiting spectral distribution of the whitened features is Marchenko-Pastur. Remarkably, while the Marchenko-Pastur analysis is far more precise near the interpolation peak, where the number of parameters is just enough to fit the training data, it coincides exactly with the distribution independent bound as the level of overparametrization increases.

1 Introduction

Classical statistics and machine learning models have traditionally been analyzed in regimes where the training loss (also known as the empirical risk) approximates the test loss. In contrast, many modern deep learning systems obtain a much lower loss on the training set than on the test set. In recent years there has been a growing theoretical understanding that different models that fit noisy data perfectly (interpolate the data) can nevertheless generalize optimally or nearly optimally. This phenomenon, which has come to be known as “benign overfitting” [1] or “harmless interpolation” [2] can be shown to provably occur in a wide array of settings including for non-parametric weighted nearest neighbor type methods [3, 4, 5, 6], linear regression [7, 1, 2, 8, 9, 10, 11, 12], random features and kernel methods [13, 14, 15, 16], and neural networks [17, 18, 19, 17, 20], to give just some representative examples of recent literature. Each of these works reveal a set of sufficient conditions on the data and learning algorithm for which benign overfitting is possible. However, a common aspect is overparametrization: the number of model parameters is significantly larger than what is required to fit all the training data.

Indeed, a striking feature of many current models is their size, reaching billions or even trillions parameters [21]. A clue to understanding the need for a large number of parameters is provided by the double descent generalization curve proposed in [22] which qualitatively describes the relation between model size and its test performance. The shape of the curve suggests that interpolating models need to be significantly over-parameterized compared to the “interpolation threshold” (the minimum number of parameters needed to fit the training data) to achieve near-optimal performance. However, despite the demonstration of double descent and benign overfitting for many specific settings, there is little literature which seeks to understand *necessary* conditions for such phenomena to occur. The most significant step in this direction was made by Holzmüller [23] which shows that double descent is universal for minimum norm (ridgeless) linear regression and implies the necessity of overparameterization for these interpolating models to have good performance.

Yet, in practice models are trained using iterative (gradient-based) methods which are typically stopped early, well before convergence to a truly interpolating solution. Indeed, pushing models to fit the data perfectly can be prohibitively computationally expensive and is usually unnecessary. Thus, while the interpolation analyses are insightful for understanding modern machine learning, they are a limit case for the settings of most practical interest.

In this paper, we aim to shed light on these issues by analyzing the connection between the *training (empirical) loss*, the *expected (test) loss* and the *number of parameters* for general linear models. Specifically, we show that for any algorithm producing a linear predictor with near-optimal expected loss, the output model must either be “classical” with the training loss relatively large – close to the noise level, or have a large excess over-parameterization (i.e., the number of parameters relative to the minimum necessary to just fit the training data without concern for generalization)¹. The trade-off between the empirical loss and the number of parameters is universal in the sense that it holds for *any algorithm* which outputs a linear predictor using *any* non-degenerate feature map for *any* regression problem with noise.

Furthermore, we provide a more precise analysis under additional distributional and asymptotic assumptions. Remarkably, the universal bound is tight (up to constant factors) for this far more special “asymptotic Gaussian” case when the model is sufficiently over-parameterized and obtains training loss strictly below the noise level.

To introduce the setting of interest, consider a general regression problem where we are given a training set of n samples $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ with $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$. We assume that each point is sampled $(\mathbf{x}_i, y_i) \sim_{i.i.d.} P$ for some distribution P on $\mathbb{R}^d \times \mathbb{R}$. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ we define its training (empirical) loss $\widehat{R}_n(f)$ and its expected (test) loss $R(f)$ as

$$\widehat{R}_n(f) := \frac{1}{n} \sum_{i=1}^n (y_i - f(\mathbf{x}_i))^2, \quad R(f) := \mathbb{E}_{(\mathbf{x}, y)} (y - f(\mathbf{x}))^2.$$

The *regression function* is defined as $f_\star(\mathbf{x}) := \mathbb{E}(y | \mathbf{x})$. It is well-known that $f_\star(\mathbf{x})$ is the optimal predictor for regression in the sense of minimizing the expected loss:

$$f_\star(\mathbf{x}) = \arg \min_f R(f).$$

Thus for an arbitrary predictor f , it makes sense to consider the *excess loss*

$$\mathcal{E}(f) := R(f) - R(f_\star)$$

as a measure of the performance of f compared to the best theoretically achievable test loss. Finally we assume that the problem has noise level of at least σ^2 , that is for almost all \mathbf{x}

$$\text{Var}(y | \mathbf{x}) \geq \sigma^2.$$

Note that the last condition directly implies that $R(f_\star) \geq \sigma^2$. We now consider a general p -dimensional linear feature model $\phi_p : \mathbb{R}^d \rightarrow \mathbb{R}^p$ of the form

$$\beta(\mathbf{x}) = \beta^\top \phi_p(\mathbf{x}), \quad \beta \in \mathbb{R}^p.$$

Here the map ϕ_p can be deterministic or random. The optimal linear predictor β_\star is given by

$$\beta_\star = \arg \min_{\beta \in \mathbb{R}^p} R(\beta).$$

It is clear that the excess loss of β compared to best predictor f_\star is bounded from below by the excess loss of β compared to the best *linear* predictor β_\star

$$\mathcal{E}(\beta) \geq \mathcal{E}^{\text{lin}}(\beta) := R(\beta) - R(\beta_\star).$$

Assume now that we have an algorithm \mathcal{A} that given the training data \mathcal{D}_n , outputs a linear predictor $\beta = \mathcal{A}(\mathcal{D}_n)$, with the empirical loss bounded by $\tau \geq 0$ relative to the noise level almost surely i.e.,

$$\widehat{R}_n(\mathcal{A}(\mathcal{D}_n)) / \sigma^2 \leq \tau.$$

Additionally, assume that ϕ_p is non-degenerate (see Section 2 for the exact conditions). Our main result is the following lower bound on the expected excess loss

$$\mathbb{E}_{\mathcal{D}_n \sim P^n} \mathcal{E}(\mathcal{A}(\mathcal{D}_n)) \geq \mathbb{E}_{\mathcal{D}_n \sim P^n} \mathcal{E}^{\text{lin}}(\mathcal{A}(\mathcal{D}_n)) \geq \begin{cases} \sigma^2 \frac{n}{p} (1 - \sqrt{\tau})^2, & \tau < 1 \\ 0 & \tau \geq 1 \end{cases} \quad (1)$$

¹The “classical” vs “modern” distinction is based on the empirical loss rather than the number of parameters. Thus, for models rich enough to fit the training data perfectly, it is a consequence of the training algorithm rather than an inherent property of the model as such. Indeed, models in classical settings can still be highly parametric, as is the case for the traditional analyses of kernel machines (e.g., [24]), which can be viewed as infinite-dimensional linear models.

Remark 1. *The case of $\tau \geq 1$ necessarily results in a trivial bound. Indeed, suppose that $\text{Var}(y | \mathbf{x}) = \sigma^2$, and the optimal predictor is linear. Then the “oracle” algorithm that always outputs β_* for any input has expected training loss $\mathbb{E}_{\mathcal{D}_n \sim P^n} [\widehat{R}_n(\beta_*)] = \sigma^2$ while its expected excess loss $\mathbb{E}_{\mathcal{D}_n \sim P^n} [\mathcal{E}(\beta_*)] = 0$. In general, our lower bound applies to any algorithm (even ones with knowledge of β_*) and for any problem instance. In contrast, minimax results lower bound the performance of any algorithm on worst case problem instances.*

While the bound above is very general, a more precise analysis is possible asymptotically under additional distributional assumptions. Assume that $n, p \rightarrow \infty$ with limiting ratio $\lim n/p = \gamma$ and that the limiting spectral distribution is Marchenko-Pastur with aspect ratio γ (see Appendix B), which is the case if for example the covariates are Gaussian. Additionally assume that the true model is linear with additive noise obeying certain mild moment conditions (see Section 3 for details). Define

$$\mathcal{E}_*(\tau) = \inf_{\mathcal{A} \text{ s.t. } \widehat{R}_n(\mathcal{A}(\mathcal{D}_n)) \leq \tau \sigma^2} \mathbb{E}_{\mathcal{D}_n \sim P^n} [\mathcal{E}(\mathcal{A}(\mathcal{D}_n))]$$

i.e., $\mathcal{E}_*(\tau)$ is the minimal expected excess loss for any algorithm with training loss at most $\tau \sigma^2$. In this setting, it turns out that the bound in Eq. (1) becomes tight when $p \gg n$ namely,

$$\mathcal{E}_*(\tau) \sim \sigma^2 \frac{n}{p} (1 - \tau^{1/2})^2 \text{ as } p/n \rightarrow \infty, \tau \in [0, 1]. \quad (2)$$

For small τ and $\gamma < 1$ the bound can be improved to the following²:

$$\mathcal{E}_*(\tau) \geq \sigma^2 \frac{n}{p-n} \left(1 - \tau^{1/2} \sqrt{\frac{p}{p-n}} \right)^2, \tau \in [0, 1 - \gamma]. \quad (3)$$

In fact, the bound is tight up to an $o(\sqrt{\tau})$ error term. Furthermore, under the same conditions, at the interpolation peak $p = n$, we have the following precise expression for the minimal expected excess loss

$$\mathcal{E}_*(\tau) = \sigma^2 \left(\frac{1}{4\tau} + \frac{\tau}{4} - \frac{1}{2} \right). \quad (4)$$

A few observations are now in order.

Comparison at interpolation ($\tau = 0$). A special case of our general result Equation (1) is for $\tau = 0$, namely when we only consider models that interpolate the data. It is instructive to compare our bound with some of the existing work in that setting. For $\tau = 0$ we obtain a lower bound for the excess loss of $\sigma^2 \cdot n/p$ which matches the result in [2] (Corollary 1). The result in [2] is given for the well-specified linear setting and requires specific covariate assumptions such as Gaussianity, but holds with high probability rather than just in expectation.

A lower bound for minimum norm interpolating linear models without distributional assumptions is given in [23]. The bound is of the form $\sigma^2 \cdot n/(p-n+1)$ and is significantly tighter near the interpolation peak $p = n$. Remarkably, their general bound which holds under minimal assumptions, almost matches the exact computation for the Gaussian case in [8] which yields $\sigma^2 \cdot n/(p-n-1)$ for $p \geq n+2$.

Remark 2. *We note that while the results in [23] are stated for minimum norm predictors, their analysis implies an algorithm independent lower bound for interpolating models, which is sharper than our bound in Equation (1) for $\tau = 0$.*

Comparison between interpolating and non-interpolating regimes. We will now compare the interpolating regime ($\tau = 0$) with the non-interpolating regime ($\tau > 0$).

a. The peak behavior ($p \approx n$). The general results in [23] demonstrate a sharp peak at the interpolation threshold $p = n$. Indeed, the analysis for the Gaussian setting [8] shows that the peak is in fact infinite. Note however, any non-zero regularization attenuates the peak, making it finite (e.g., [14]). Note that τ can also be viewed as regularization. In the asymptotic MP setting Equation (4) shows that the expected loss has a pole singularity τ^{-1}

²Note that in the limit $\frac{n}{p-n} \sim \frac{\gamma}{1-\gamma}$ and $\frac{p}{p-n} \sim \frac{1}{1-\gamma}$.

at the peak. Thus the transition between interpolating and non-interpolating regimes is discontinuous in terms of the height of the interpolation peak. Hence we see that our general bound in Equation (1) is loose close to the peak, which is to be expected as it is continuous in τ at 0, while the actual expected loss is discontinuous. In contrast, the Marchenko-Pastur setting bound in Equation (3) is much more accurate.

b. The “tail” behavior ($p \gg n$). In sharp contrast to the peak, our general bound show that the “tail” ($p \gg n$) behavior of the generalization curve is remarkably stable with respect to τ . Achieving nearly optimal excess risk requires either $\tau > 1 - o(1)$ or $p \gg n$ (or both, of course). Thus interpolating and non-interpolating solutions require essentially the same level of over-parameterization to approach optimality as long as the loss of non-interpolating models is at least slightly lower than the noise level.

Furthermore, for any empirical loss smaller than the noise level ($\tau < 1$), when $p \gg n$ the general bound in Equation (1) matches the asymptotic analysis in the Marchenko-Pastur setting. This is remarkable, as the general bound is not asymptotic and makes essentially no assumptions on the covariate distribution, regression function, or the structure of the noise, yet it is still tight in the limit of increasing overparametrization.

While the interpolation peak is a striking feature of the generalization curve, the tail behavior is arguably more important for understanding practical applications. Indeed, the tail behaviour seems consistent with over-parameterization in practical models which are routinely trained so that the training loss is significantly lower than the test loss but far from zero. In contrast, the peak is a less robust phenomenon which describes only specific regime of training and is highly sensitive to the presence of regularization.

Convergence rates. A classical line of statistical analysis is concerned with convergence rates for various estimation problems. Typically statistical rates for regression (see e.g., [25]) are of the form

$$\mathcal{E}(\hat{f}_n) = O(n^{-\alpha}), \quad \alpha > 0$$

where \hat{f}_n is a predictor based on a training set with n samples and α may depend on some notion of data dimensionality, such as the dimension of the data manifold.

We note that such parametric or non-parametric rates are easily compatible with our analysis. As a corollary of our lower bound Equation (1), we can see that in order to achieve such a rate with a linear model the model either needs to have excess over-parameterization inversely proportional to the rate, i.e., $p = \Omega(n^{1+\alpha})$ or to be in the “classical regime” where $\tau = 1 - o_n(1)$ i.e., the train error essentially at the noise level. Note that in general (e.g., for a random feature model [26]) the number of features p is a property of the model and is distinct from any notion of data dimensionality. Moreover, note that to simply achieve low training error it is only necessary that $p = \Omega(n)$ as only $p = n$ features are necessary to interpolate the training data, hence the condition $p = \Omega(n^{1+\alpha})$ is additionally requiring at least n^α times the number of parameters to achieve the desired excess loss rate.

Interestingly, there are also settings where near-interpolation is necessary to approach optimal generalization [27, 28, 29]. Our results imply that in these settings significant excess over-parameterization is unavoidable. Of the aforementioned works, the most closely related to ours work [27] which studies high-dimensional linear regression. Moreover, they study the optimal test loss subject to a training loss constraint as we do in this paper. However, a key difference is that their results hold in a Bayesian setting where the true model is drawn from some prior distribution and all losses are averaged over this prior whereas our results hold even for a fixed target function. In particular, only in our setting can an estimator achieve zero excess loss using finitely many samples.

Finally, we note that the trade-off presented in this paper is reminiscent of the trade-off between smoothness and over-parameterization discovered in [30], which shows that over-parameterization is necessary to fit noisy data smoothly. In contrast to [30] which does not consider generalization, predictors that generalizes well while fitting noise need to be “spiky” rather than smooth, hence over-parameterization in our paper serves a different function.

2 Universal Lower Bound

First let us introduce some notation. We define the training data matrices

$$\mathbf{X} := (\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top) \in \mathbb{R}^{n \times d}, \quad \mathbf{y} := (y_1, \dots, y_n) \in \mathbb{R}^n$$

so that $\mathcal{D}_n = (\mathbf{X}, \mathbf{y}) \sim P^n$. If $(\mathbf{x}, y) \sim P$ then we will use P_X to denote the marginal distribution on \mathbf{x} . We will use the following abbreviated notation for the expectations

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} := \mathbb{E}_{\mathcal{D}_n \sim P^n}, \quad \mathbb{E}_{(\mathbf{x}, y)} := \mathbb{E}_{(\mathbf{x}, y) \sim P}, \quad \mathbb{E}_{\mathbf{x}} := \mathbb{E}_{\mathbf{x} \sim P_X}.$$

We define the feature matrix and feature covariance matrix

$$\Phi := (\phi_p(\mathbf{x}_1)^\top, \dots, \phi_p(\mathbf{x}_n)^\top) \in \mathbb{R}^{n \times p}, \quad \Sigma := \mathbb{E}_{\mathbf{x}}[\phi_p(\mathbf{x})\phi_p(\mathbf{x})^\top] \in \mathbb{R}^{p \times p}.$$

We will also make use of the whitened feature-matrix

$$\mathbf{W} := \Phi \Sigma^{-1/2} \in \mathbb{R}^{n \times p}$$

and the whitened empirical covariance matrix

$$\mathbf{G} := \frac{1}{p} \mathbf{W} \mathbf{W}^\top \in \mathbb{R}^{n \times n}.$$

We now state our main result

Theorem 1 (Universal Lower Bound). *Let $n, p \geq 1$. Assume that P and ϕ_p satisfy the following*

1. $\mathbb{E}y^2 < \infty$ which implies $R(f_\star) < \infty$,
2. $\text{Var}(y | \mathbf{x}) \geq \sigma^2$ almost surely over \mathbf{x} ,
3. $\text{rank}(\Phi) = \min(n, p)$ almost surely.

Then for any algorithm \mathcal{A} which outputs a linear feature model $\mathcal{A}(\mathcal{D}_n)$ with training loss almost surely at most $\tau\sigma^2$

$$\mathbb{E}_{\mathcal{D}_n \sim P^n} \mathcal{E}(\mathcal{A}(\mathcal{D}_n)) \geq \sigma^2 \frac{n}{p} (1 - \sqrt{\tau})^2.$$

Remark 3. *Note that we can consider the feature map ϕ_p to be random as well, for instance taking ϕ_p to be a neural network with output dimension p and random weights θ . It is often the case that Assumption 3 will hold almost surely over θ (see Theorem 10 in [23]). Since the weights are independent of the data, as an immediate corollary to Theorem 1 we get the same lower bound when additionally taking expectation over θ .*

Now that we have stated our main result and gave some of its interpretations and consequences, we will move on to giving its proof which is pleasantly elementary. We will start by setting up relevant definitions and providing some starting lemmas, before moving on to the core proof.

2.1 Proof of Theorem 1

Under the assumptions of Theorem 1 we define the minimal excess test loss of τ -overfitting p -dimensional linear feature models trained on the dataset \mathcal{D}_n of n samples as

$$\mathcal{E}_\star(\tau; n, p) := \min_{\beta \in \mathbb{R}^p} \mathcal{E}(\beta) \text{ s.t. } \widehat{R}_n(\beta) \leq \tau\sigma^2.$$

It then suffices to show that

$$\mathbb{E}_{\mathcal{D}_n \sim P^n} [\mathcal{E}_\star(\tau; n, p)] \geq \sigma^2 \frac{n}{p} (1 - \sqrt{\tau})^2.$$

We will bound the excess test loss $\mathcal{E}(\beta)$ by comparing it with the excess loss with respect to the optimal linear model $\mathcal{E}^{\text{lin}}(\beta)$. Let us denote the optimal linear predictor as

$$\beta_\star := \arg \min_{\beta \in \mathbb{R}^p} R(\beta). \tag{5}$$

For further analysis we will need the following two basic lemmas which are standard results characterizing the excess linear loss. The proofs can be found in Appendix A.

Lemma 1 (Optimal Linear Predictor). *Define β_* as in Eq. (5). Then*

1. β_* is the orthogonal projection of f_* onto the subspace of linear functions $\mathcal{H} = \{\beta(\mathbf{x}) : \beta \in \mathbb{R}^p\}$ in L^2 ,
2. $\mathcal{E}(f) := R(f) - R(f_*) = \mathbb{E}_{\mathbf{x}}[(f(\mathbf{x}) - f_*(\mathbf{x}))^2]$,
3. $\mathcal{E}^{\text{lin}}(\beta) := R(\beta) - R(\beta_*) = \mathbb{E}_{\mathbf{x}}[(\beta(\mathbf{x}) - \beta_*(\mathbf{x}))^2]$.

The first claim above establishes the equivalent characterization of the optimal linear predictor as a projection of the optimal predictor onto the space of linear functions. The second claim asserts that the excess loss of a given function is its L^2 distance to the optimal function, and similarly the third claim asserts that the excess loss a linear function is its L^2 distance to the optimal linear function. We now state the second lemma.

Lemma 2 (Excess Linear Loss). *The excess loss is lower bounded by the excess linear loss, that is $\mathcal{E}(\beta) \geq \mathcal{E}^{\text{lin}}(\beta)$. Moreover we can write the excess linear loss explicitly as*

$$\mathcal{E}^{\text{lin}}(\beta) = \left\| \Sigma^{1/2}(\beta - \beta_*) \right\|^2.$$

The results above concern the excess (linear) loss of a given predictor. However to establish lower bounds we will consider the minimal value of this quantity subject to a training loss constraint. Define the minimal excess linear loss $\mathcal{E}_*^{\text{lin}}(\tau)$ for training dataset \mathcal{D}_n as

$$\mathcal{E}_*^{\text{lin}}(\tau) = \min_{\beta \in \mathbb{R}^p} \mathcal{E}^{\text{lin}}(\beta) \text{ s.t. } \widehat{R}_n(\beta) \leq \tau\sigma^2.$$

Let $\mathbf{W} = \Phi \Sigma^{-1/2}$ be the whitened features. Define the random vectors

$$\xi_{\text{noise}} = (y_i - f_*(\mathbf{x}_i))_{i \in [n]} \in \mathbb{R}^n, \quad \xi_{\text{lin}} = (f_*(\mathbf{x}_i) - \beta_*(\mathbf{x}_i))_{i \in [n]} \in \mathbb{R}^n$$

and let $\xi = \xi_{\text{noise}} + \xi_{\text{lin}}$. We give an alternate optimization problem for characterizing the minimal excess linear loss which will be more amenable to analysis later on. The same equivalence (for $\tau = 0$) appears in the proof of Theorem 1 in [2], however as it is not a very standard result in the literature and is crucial to the rest of our analysis, we record the statement and its proof here.

Lemma 3 (Minimal Excess Linear Loss). *We can equivalently write $\mathcal{E}_*^{\text{lin}}(\tau)$ as*

$$\mathcal{E}_*^{\text{lin}}(\tau) = \min_{\mathbf{b} \in \mathbb{R}^p} \|\mathbf{b}\|_2^2 \text{ s.t. } \frac{1}{n} \|\mathbf{W}\mathbf{b} - \xi\|_2^2 \leq \tau\sigma^2, \quad (6)$$

and $\mathcal{E}_*(\tau) \geq \mathcal{E}_*^{\text{lin}}(\tau)$.

Proof. By definition we can write the training loss as

$$\widehat{R}_n(\beta) = \frac{1}{n} \|\Phi\beta - \mathbf{y}\|^2.$$

Note that $\mathbf{y} = \Phi\beta_* + \xi$, hence

$$\begin{aligned} \Phi\beta - \mathbf{y} &= \Phi\beta - (\Phi\beta_* + \xi) \\ &= \Phi(\beta - \beta_*) - \xi \\ &= \Phi \Sigma^{-1/2} \Sigma^{1/2}(\beta - \beta_*) - \xi \\ &= \mathbf{W} \Sigma^{1/2}(\beta - \beta_*) - \xi. \end{aligned}$$

Therefore by Lemma 2 we see that $\mathcal{E}_*^{\text{lin}}(\tau)$ is equivalent to

$$\min_{\beta} \left\| \Sigma^{1/2}(\beta - \beta_*) \right\|^2 \text{ s.t. } \frac{1}{n} \left\| \mathbf{W} \Sigma^{1/2}(\beta - \beta_*) - \xi \right\|^2 \leq \tau\sigma^2,$$

which after making the following substitution over the optimization variable

$$\mathbf{b} = \Sigma^{1/2}(\beta - \beta_*)$$

can be seen to be equivalent to Eq. (6). □

Before proceeding to the proof of the theorem, we state one more lemma which will be useful for dealing with quadratic forms involving the noise vector. The proof is in Appendix A.

Lemma 4 (Expectation Over Noise). *Let $f : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}_+^n$ be any PSD matrix valued function. Then,*

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}}[\boldsymbol{\xi}^\top f(\mathbf{X}) \boldsymbol{\xi}] \geq \sigma^2 \mathbb{E}_{\mathbf{X}} \text{Tr}(f(\mathbf{X})).$$

Moreover, if $\text{Var}(y | \mathbf{x}) = \sigma^2$ almost surely and $f_\star = \beta_\star$ then the above is an equality.

We are now ready to give the proof of our main Theorem 1. We will proceed to lower bound $\mathcal{E}_\star^{\text{lin}}(\tau)$ which will then imply the lower bound in Theorem 1 by the previous lemmas. To lower bound $\mathcal{E}_\star^{\text{lin}}(\tau)$ we will apply weak duality to the optimization problem in Lemma 3. Interestingly, as we will see later (see Eq. (13)), it turns out that this dual lower bound reveals that the minimal excess linear loss can be bounded below by the test loss of a ridge regression estimator on an auxiliary problem coming from Lemma 3 (see Remark 5).

Proof of Theorem 1. Consider the optimization problem in Eq. (6) defining $\mathcal{E}_\star^{\text{lin}}(\tau)$. The associated Lagrangian is

$$\mathcal{L}(\mathbf{b}, \eta) = \|\mathbf{b}\|_2^2 + \eta \left(\frac{1}{n} \|\mathbf{W}\mathbf{b} - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right)$$

for $\eta \geq 0$ and the dual function is

$$g(\eta) = \inf_{\mathbf{b}} \mathcal{L}(\mathbf{b}, \eta).$$

After some rearrangement, we can rewrite the Lagrangian as

$$\mathcal{L}(\mathbf{b}, \eta) = \mathbf{b}^\top \left(\mathbf{I} + \frac{\eta}{n} \mathbf{W}^\top \mathbf{W} \right) \mathbf{b} - \frac{2\eta}{n} \mathbf{b}^\top \mathbf{W} \boldsymbol{\xi} + \frac{\eta}{n} \|\boldsymbol{\xi}\|_2^2 - \eta\tau\sigma^2$$

which is a convex quadratic objective in \mathbf{b} . Hence we can minimize it by setting the derivative to zero. The gradient of $\mathcal{L}(\mathbf{b}, \eta)$ is given by

$$\nabla_{\mathbf{b}} \mathcal{L}(\mathbf{b}, \eta) = 2 \left(\mathbf{I} + \frac{\eta}{n} \mathbf{W}^\top \mathbf{W} \right) \mathbf{b} - \frac{2\eta}{n} \mathbf{W}^\top \boldsymbol{\xi}$$

hence setting this equal to zero and solving for the optimal \mathbf{b} we get

$$\begin{aligned} \mathbf{b}_\star(\eta) &= \arg \min_{\mathbf{b}} \mathcal{L}(\mathbf{b}, \eta) \\ &= \eta(n\mathbf{I} + \eta\mathbf{W}^\top \mathbf{W})^{-1} \mathbf{W}^\top \boldsymbol{\xi} \\ &= \mathbf{W}^\top ((n/\eta)\mathbf{I} + \mathbf{W}\mathbf{W}^\top)^{-1} \boldsymbol{\xi} \\ &= \frac{1}{p} \mathbf{W}^\top \left(\frac{n}{p} \frac{1}{\eta} \mathbf{I} + \frac{1}{p} \mathbf{W}\mathbf{W}^\top \right)^{-1} \boldsymbol{\xi}. \end{aligned}$$

By weak duality we have that for any $\eta \geq 0$,

$$\mathcal{E}_\star^{\text{lin}}(\tau) \geq g(\eta) = \mathcal{L}(\mathbf{b}_\star(\eta), \eta).$$

For convenience we make the following change of variables

$$\lambda := \frac{n}{p} \frac{1}{\eta}$$

which is a bijection between $[0, +\infty]$ and itself. Thus in terms of λ we can write

$$\mathcal{L}(\mathbf{b}, \lambda) = \|\mathbf{b}\|_2^2 + \frac{n}{p\lambda} \left(\frac{1}{n} \|\mathbf{W}\mathbf{b} - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right) \tag{7}$$

$$\mathbf{b}_\star(\lambda) = \frac{1}{p} \mathbf{W}^\top \left(\lambda \mathbf{I} + \frac{1}{p} \mathbf{W}\mathbf{W}^\top \right)^{-1} \boldsymbol{\xi} \tag{8}$$

where interestingly $\mathbf{b}_*(\lambda)$ happens to be the ridge regression estimator with ridge parameter λ on whitened covariates \mathbf{W} with pure noise target $\boldsymbol{\xi}$. For all $\lambda \geq 0$ we have the lower bound

$$\mathcal{E}_*^{\text{lin}}(\tau) \geq g(\lambda).$$

Taking expectations, we have the following bound

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} [\mathcal{E}_*^{\text{lin}}(\tau)] \geq \mathbb{E}_{\mathbf{y}, \mathbf{X}} [\mathcal{L}(\mathbf{b}_*(\lambda), \lambda)] = \mathbb{E}_{\mathbf{y}, \mathbf{X}} \|\mathbf{b}_*(\lambda)\|^2 + \frac{n}{p\lambda} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \left(\frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda) - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right). \quad (9)$$

Let $\mathbf{G} = \mathbf{W}\mathbf{W}^\top/p$ and $\mathbf{G}(\lambda) = \mathbf{G} + \lambda\mathbf{I}$. Denote the eigenvalues of \mathbf{G} as $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. We will show

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} \|\mathbf{b}_*(\lambda)\|^2 \geq \frac{\sigma^2}{p} \mathbb{E}_{\mathbf{X}} \text{Tr}(\mathbf{G}(\lambda)^{-1} \mathbf{G} \mathbf{G}(\lambda)^{-1}) = \frac{\sigma^2}{p} \mathbb{E}_{\mathbf{X}} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda)^2} \quad (10)$$

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} \frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda) - \boldsymbol{\xi}\|_2^2 \geq \frac{\sigma^2}{n} \mathbb{E}_{\mathbf{X}} \text{Tr}(\mathbf{G} \mathbf{G}(\lambda)^{-1} - \mathbf{I})^2 = \frac{\sigma^2}{n} \mathbb{E}_{\mathbf{X}} \sum_{i=1}^n \left(\frac{\lambda}{\lambda_i + \lambda} \right)^2 \quad (11)$$

using Lemma 4. For Eq. (10) we have the following

$$\begin{aligned} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \|\mathbf{b}_*(\lambda)\|^2 &= \frac{1}{p} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \left[\boldsymbol{\xi}^\top \mathbf{G}(\lambda)^{-1} \mathbf{G} \mathbf{G}(\lambda)^{-1} \boldsymbol{\xi} \right] \\ &\geq \frac{\sigma^2}{p} \mathbb{E}_{\mathbf{X}} \text{Tr}(\mathbf{G}(\lambda)^{-1} \mathbf{G} \mathbf{G}(\lambda)^{-1}). \end{aligned}$$

Similarly, for Eq. (11)

$$\begin{aligned} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda) - \boldsymbol{\xi}\|_2^2 &= \frac{1}{n} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \left[\boldsymbol{\xi}^\top (\mathbf{G} \mathbf{G}(\lambda)^{-1} - \mathbf{I})^2 \boldsymbol{\xi} \right] \\ &\geq \frac{\sigma^2}{n} \mathbb{E}_{\mathbf{X}} \text{Tr}(\mathbf{G} \mathbf{G}(\lambda)^{-1} - \mathbf{I})^2. \end{aligned}$$

Note that by the full rank assumption, $\lambda_1 \geq \dots \geq \lambda_{\min(n,p)} > 0$ and $\lambda_{\min(n,p)+1} = \dots = \lambda_n = 0$ almost surely. Define the function $f : [0, +\infty) \rightarrow \mathbb{R}$ as

$$f(\lambda) := \frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda}{\lambda_i + \lambda} \right)^2 = \frac{1}{n} \sum_{i=1}^{\min(n,p)} \left(\frac{\lambda}{\lambda_i + \lambda} \right)^2 + \max(0, 1 - p/n).$$

Note that f is continuous and $\inf_{\lambda} f(\lambda) = f(0) = \max(0, 1 - p/n)$ and $\sup_{\lambda} f(\lambda) = f(\infty) = 1$. Therefore as long as $1 - p/n \leq \tau$ there exists λ_* such that

$$f(\lambda_*) = \frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda_*}{\lambda_i + \lambda_*} \right)^2 = \tau. \quad (12)$$

Otherwise if $1 - p/n > \tau$, then by taking $\lambda \rightarrow 0$, we have that

$$\frac{1}{\lambda} \cdot \mathbb{E}_{\mathbf{y}, \mathbf{X}} \left(\frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda) - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right) \rightarrow \infty,$$

which by inequality Eq. (9) implies $\mathbb{E}_{\mathbf{y}, \mathbf{X}} [\mathcal{E}_*^{\text{lin}}(\tau)] = +\infty$ and any lower bound holds trivially. Thus we will assume that $p/n \geq 1 - \tau$, in which case we will be able to obtain a non-vacuous result. Let λ_* be the random variable dependent on \mathbf{X} that satisfies Eq. (12). Then by Eq. (11) we have that

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} \left(\frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda_*) - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right) \geq 0.$$

Thus from Eq. (9) we have that

$$\begin{aligned}\mathbb{E}_{\mathbf{y}, \mathbf{X}} [\mathcal{E}_*^{\text{lin}}(\tau)] &\geq \mathbb{E}_{\mathbf{y}, \mathbf{X}} \|\mathbf{b}_*(\lambda_*)\|^2 + \frac{n}{p\lambda_*} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \left(\frac{1}{n} \|\mathbf{W}\mathbf{b}_*(\lambda_*) - \boldsymbol{\xi}\|_2^2 - \tau\sigma^2 \right) \\ &\geq \mathbb{E}_{\mathbf{y}, \mathbf{X}} \|\mathbf{b}_*(\lambda_*)\|^2.\end{aligned}\quad (13)$$

Hence from Eq. (10) we have

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}} [\mathcal{E}_*^{\text{lin}}(\tau)] \geq \sigma^2 \frac{n}{p} \cdot \mathbb{E}_{\mathbf{X}} \frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2}.\quad (14)$$

Thus to lower bound $\mathbb{E}[\mathcal{E}_*^{\text{lin}}(\tau)]$ we can try to lower bound the following

$$\frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} \quad \text{subject to} \quad \frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda_i}{\lambda_i + \lambda_*} \right)^2 = \tau,\quad (15)$$

by a quantity that we can later easily bound in expectation over \mathbf{X} . Note that if $\lambda_1 = \dots = \lambda_n = \lambda$, then

$$\lambda_* = \lambda \frac{\sqrt{\tau}}{(1 - \sqrt{\tau})} \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} = \frac{1}{\lambda} (1 - \sqrt{\tau})^2\quad (16)$$

and if $\lambda_* = 0$ then

$$\sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} = \frac{1}{n} \sum_{i=1}^n \frac{1}{\lambda_i} \geq \frac{n}{\sum_{i=1}^n \lambda_i}\quad (17)$$

where the last inequality is the AM-HM inequality (Lemma 11). Therefore a possible lower bound for Eq. (15) is

$$\frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} \geq \frac{n}{\sum_{i=1}^n \lambda_i} (1 - \sqrt{\tau})^2,\quad (18)$$

as it satisfies the edge cases in Eqs. (16) and (17). We will show that this inequality in fact holds.

By Eq. (12), to prove that Eq. (18) holds, it suffices to show that for any $a_1 \geq \dots \geq a_n \geq 0$ and $x \geq 0$,

$$\frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^2} \geq \frac{n}{\sum_{i=1}^n a_i} \left(1 - \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{x}{a_i + x} \right)^2} \right)^2$$

which upon rearranging is equivalent to

$$\sqrt{\frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^2}} \sqrt{\frac{1}{n} \sum_{i=1}^n a_i} + \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{x}{a_i + x} \right)^2} \geq 1.\quad (19)$$

Let us fix $a_1 \geq \dots \geq a_n \geq 0$ and consider the function $g : [0, \infty) \rightarrow [0, \infty)$ defined as the left-hand side of the above

$$g(x) = \sqrt{\frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^2}} \sqrt{\frac{1}{n} \sum_{i=1}^n a_i} + \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{x}{a_i + x} \right)^2}.\quad (20)$$

We can show $g(x)$ is decreasing by computing the derivative which is given as follows

$$g'(x) = \frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^3} \cdot \left[\left(\frac{1}{n} \sum_{i=1}^n \frac{1}{(a_i + x)^2} \right)^{-1/2} - \left(\frac{1}{n} \sum_{i=1}^n a_i \right)^{1/2} \cdot \left(\frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^2} \right)^{-1/2} \right].$$

Therefore to show that $g'(x) \leq 0$, by rearranging the above expression it suffices show that

$$\frac{1}{n} \sum_{i=1}^n \frac{a_i}{(a_i + x)^2} \leq \left(\frac{1}{n} \sum_{i=1}^n a_i \right) \cdot \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{(x + a_i)^2} \right).$$

The above however follows immediately from Chebyshev's Sum Inequality (Lemma 12). Therefore we have shown that $g(x)$ is decreasing. Since in Eq. (20) it is easy to see that $\lim_{x \rightarrow \infty} g(x) = 1$, it follows that $g(x) \geq 1$ for all $x \geq 0$, which proves Eq. (19). Now taking Eq. (19) and plugging in to Eq. (14) we get that

$$\begin{aligned} \mathbb{E}_{\mathbf{y}, \mathbf{X}} \mathcal{E}_*^{\text{lin}}(\tau) &\geq \sigma^2 \frac{n}{p} \cdot \mathbb{E}_{\mathbf{X}} \frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} \\ &\geq \sigma^2 \frac{n}{p} \cdot \mathbb{E}_{\mathbf{X}} \frac{n}{\sum_{i=1}^n \lambda_i} (1 - \sqrt{\tau})^2 \\ &\geq \sigma^2 \frac{n}{p} (1 - \sqrt{\tau})^2 \frac{n}{\mathbb{E}_{\mathbf{X}} \sum_{i=1}^n \lambda_i} && \text{(Jensen's Inequality)} \\ &= \sigma^2 \frac{n}{p} (1 - \sqrt{\tau})^2, \end{aligned}$$

where the last equality holds since $\mathbb{E}_{\mathbf{X}} \sum_{i=1}^n \lambda_i = \mathbb{E}_{\mathbf{X}} \text{Tr}(\mathbf{G}) = n$. □

Remark 4. One may wonder if the following lower bound

$$\frac{1}{n} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda_*)^2} \geq \frac{1}{n} \sum_{i=1}^n \frac{1}{\lambda_i} (1 - \sqrt{\tau})^2, \quad (21)$$

could have been used in place of Eq. (18) as it also satisfies the same edge cases. This however is not a valid inequality. To see this, take $n = 2$ and let $\lambda_1 = \varepsilon$, $\lambda_2 = 1 - \varepsilon$, and $\tau \in (1/2, 1)$. Then as $\varepsilon \rightarrow 0$, it is easy to see since

$$\tau = \frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda_i}{\lambda_i + \lambda_*} \right)^2 \approx \frac{1}{2} + \frac{1}{2} \left(\frac{\lambda_*}{1 + \lambda_*} \right)^2$$

it must be that λ_* is bounded below and does not go to 0. However that means that the left-hand side of Eq. (21) is bounded above whereas the right hand side goes to infinity.

Remark 5. The lower bound via weak duality (see Eq. (13)) reveals that the minimal excess linear loss is lower bounded by the test loss of a ridge estimator on a different, auxiliary problem coming from Lemma 3. In this auxiliary problem the covariates are whitened and the training targets are the noise $\boldsymbol{\xi}$ as opposed to \mathbf{y} . The ridge parameter of the estimator is exactly the Lagrange multiplier corresponding to the training error constraint. This ridge estimator requires oracle knowledge of the feature covariance $\boldsymbol{\Sigma}$ and the noise vector $\boldsymbol{\xi}$. Note that the lower bound arises purely from the variance due to label noise since this oracle ridge estimator is an unbiased estimator for the original problem.

3 Lower Bounds under Marchenko-Pastur Asymptotics

In this section we analyze $\mathcal{E}_*(\tau; n, p)$ more precisely under additional asymptotic distributional assumptions. As mentioned earlier, this will allow us to assess the tightness of our general lower bound given in Theorem 1. Specifically, we assume the following setting which has been used to analyse ridge regression in several prior works including [31, 32, 7]:

[MP] Assume $n, p \rightarrow \infty$ so that $n/p \rightarrow \gamma \in (0, +\infty)$. Recall the matrix $\mathbf{G} = \mathbf{W}\mathbf{W}^\top/p$ which has eigenvalues $\lambda_1, \dots, \lambda_n$. Let μ_p be the empirical spectral distribution of \mathbf{G} ,

$$\mu_p(A) = \frac{1}{n} |\{\lambda_i \in A\}|, \quad A \subset \mathbb{R}.$$

Then we have the following convergence in the weak topology

$$\mu_p \rightarrow \text{MP}(\gamma)$$

where $\text{MP}(\gamma)$ is the Marchenko-Pastur distribution with aspect ratio γ (see Appendix B).

[Lin] The optimal model is linear and noise is additive:

$$y = \beta_\star^\top \phi(\mathbf{x}) + \varepsilon,$$

where the noise ε satisfies $\mathbb{E}[\varepsilon] = 0$, $\mathbb{E}[\varepsilon^2] = \sigma^2$, and $\mathbb{E}[\varepsilon^{4+\eta}] < \infty$ for some $\eta > 0$.

Remark 6. Assumption **MP** holds whenever the entries of \mathbf{W} are distributed i.i.d with mean zero and variance 1. In particular, this holds if $\phi(\mathbf{x}) \sim \mathcal{N}(0, \Sigma)$ for any $\Sigma \succ \mathbf{0}$. By assuming the optimal model is linear we have that $\mathcal{E}_\star(\tau) = \mathcal{E}_\star^{\text{lin}}(\tau)$. Additionally, since $\text{Var}(y | \mathbf{x}) = \text{Var}(\varepsilon) = \sigma^2$ almost surely, the inequality in Lemma 4 is an equality.

Let $m(z; \gamma)$ denote the Stieltjes transform of the Marchenko-Pastur Law and let $m'(z; \gamma)$ be the derivative with respect to z (see Appendix B). Under the above assumptions we have the following analytical characterization of the asymptotic minimal excess error.

Proposition 1 (MP Asymptotics). *Under Assumptions **MP** and **Lin**, the asymptotic minimal excess error is given by the following analytical expression*

$$\mathcal{E}_\star(\tau, \gamma) := \lim_{\substack{n, p \rightarrow \infty \\ n/p \rightarrow \gamma}} \mathcal{E}_\star(\tau; n, p) = \sigma^2 \gamma [m(-\lambda; \gamma) - \lambda m'(-\lambda; \gamma)] \quad (22)$$

where λ satisfies the fixed point equation

$$\lambda^2 m'(-\lambda; \gamma) = \tau. \quad (23)$$

Proof. For any $\lambda \geq 0$, by Assumption **Lin** we can use Proposition 13 to get the asymptotic versions of Eqs. (10), (11)

$$\begin{aligned} \|\mathbf{b}_\star(\lambda)\|^2 &\sim \sigma^2 \frac{1}{p} \text{Tr}(\mathbf{G}(\lambda)^{-1} \mathbf{G} \mathbf{G}(\lambda)^{-1}) \\ \frac{1}{n} \|\mathbf{W} \mathbf{b}_\star(\lambda) - \boldsymbol{\xi}\|^2 &\sim \sigma^2 \frac{1}{n} \text{Tr}(\mathbf{G} \mathbf{G}(\lambda)^{-1} - \mathbf{I})^2 \end{aligned}$$

where $a_n \sim b_n$ denotes that $|a_n - b_n| \rightarrow 0$ almost surely. Then by Assumption **MP**

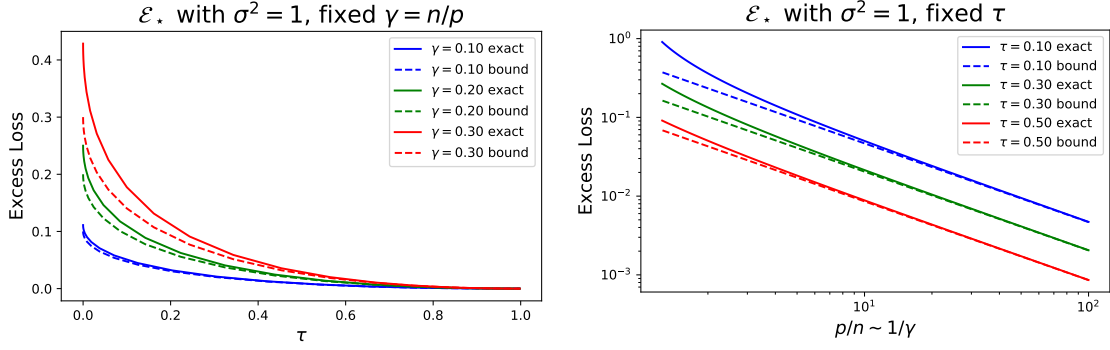
$$\begin{aligned} \frac{1}{p} \text{Tr}(\mathbf{G}(\lambda)^{-1} \mathbf{G} \mathbf{G}(\lambda)^{-1}) &= \frac{1}{p} \sum_{i=1}^n \frac{\lambda_i}{(\lambda_i + \lambda)^2} = \frac{n}{p} \int \frac{s}{(s + \lambda)^2} d\mu_p(s) \rightarrow \gamma \int \frac{s}{(s + \lambda)^2} d\text{MP}_\gamma(s), \\ \frac{1}{n} \text{Tr}(\mathbf{G} \mathbf{G}(\lambda)^{-1} - \mathbf{I})^2 &= \frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda}{\lambda_i + \lambda} \right)^2 = \int \left(\frac{\lambda}{s + \lambda} \right)^2 d\mu_p(s) \rightarrow \int \left(\frac{\lambda}{s + \lambda} \right)^2 d\text{MP}_\gamma(s). \end{aligned}$$

Therefore if λ_\star is the unique solution to the fixed point equation

$$\int \left(\frac{\lambda_\star}{s + \lambda_\star} \right)^2 d\text{MP}_\gamma(s) = \tau \quad (24)$$

then the constraint becomes tight almost surely

$$\frac{1}{n} \|\mathbf{W} \mathbf{b}_\star(\lambda_\star) - \boldsymbol{\xi}\|^2 - \tau \sigma^2 \rightarrow 0.$$



(a) For fixed $\gamma = n/p$ we plot the exact value of $\mathcal{E}_*(\tau)$ using Eq. (22) [solid] and the lower bound from Eq. (28) [dashed]. (b) For fixed τ we plot the exact value of $\mathcal{E}_*(1/\gamma)$ using Eq. (22) [solid] and the lower bound from Eq. (28) [dashed].

The pair $(\mathbf{b}(\lambda_*), \lambda_*)$ is the unique KKT point and since a minimizer of the primal problem exists due to continuity of the objective and compactness of the constraint set, this pair is asymptotically the primal/dual optimal variables and

$$\mathcal{E}_*(\tau; n, p) \sim \|\mathbf{b}_*(\lambda_*)\|^2 \rightarrow \sigma^2 \gamma \int \frac{s}{(s + \lambda_*)^2} d\text{MP}_\gamma(s)$$

where λ_* satisfies Eq. (24). We can write the integrals that appear above in terms of the Stieltjes transform

$$\begin{aligned} \int \frac{s}{(s+z)^2} d\text{MP}_\gamma(s) &= \int \left[\frac{1}{s+z} - \frac{z}{(s+z)^2} \right] d\text{MP}_\gamma(s) = m(-z; \gamma) - z m'(-z; \gamma), \\ \int \left(\frac{z}{s+z} \right)^2 d\text{MP}_\gamma(s) &= z^2 \int \frac{1}{(s+z)^2} d\text{MP}_\gamma(s) = z^2 m'(-z; \gamma). \end{aligned}$$

Thus we can write $\mathcal{E}_*(\tau, \gamma)$ as

$$\sigma^2 \gamma [m(-\lambda; \gamma) - \lambda m'(-\lambda; \gamma)] \text{ s.t. } \lambda^2 m'(-\lambda; \gamma) = \tau.$$

□

Remark 7. Note that by Remark 5, $\mathcal{E}_*(\tau, \gamma)$ should be equal to the asymptotic test risk of a ridge regression predictor on isotropic Gaussian covariates with ridge parameter λ satisfying Eq. (23), when the target function is zero. Indeed our calculations match results obtained in previous calculations of this limiting risk, for example Corollary 5 in [7].

For convenience in later proofs, let us define the following functions

$$E(\lambda, \gamma) = m(-\lambda; \gamma) - \lambda m'(-\lambda; \gamma), \quad (25)$$

$$f(\lambda, \gamma) = \lambda^2 m'(-\lambda; \gamma). \quad (26)$$

Note that f is strictly increasing in λ , so we can define the inverse function $f^{-1}(\tau, \gamma)$ so that

$$f(f^{-1}(\tau, \gamma), \gamma) = \tau. \quad (27)$$

Using our characterization of the minimal excess error in Proposition 1 we will now derive lower bounds in Theorems 2 and 4 and an exact expression when $\gamma = 1$ in Theorem 5.

Theorem 2 (MP Lower Bound). By Theorem 1, for fixed $\tau \in [0, 1]$, the minimum excess error satisfies

$$\mathcal{E}_*(\tau, \gamma) \geq \sigma^2 \gamma (1 - \sqrt{\tau})^2 \quad (28)$$

for all $\gamma \in (0, +\infty)$. The dependence on τ is in fact tight since

$$\inf_{\gamma \in (0, \infty)} \mathcal{E}_*(\tau, \gamma) / \gamma = \sigma^2 (1 - \sqrt{\tau})^2. \quad (29)$$

Proof. Observe that as $\gamma \rightarrow 0$, the spectral eigenvalue distribution of $\mathbf{W}\mathbf{W}^\top/p$ is converging to a point mass at 1. Therefore, recalling the function f defined in Eq. (26), by the Dominated Convergence Theorem

$$\lim_{\gamma \rightarrow 0} f(\lambda, \gamma) = \lim_{\gamma \rightarrow 0} \int_{s \in [1-\sqrt{\gamma}, 1+\sqrt{\gamma}]} \frac{\lambda^2}{(s+\lambda)^2} dH_\gamma(s) = \frac{\lambda^2}{(1+\lambda)^2}$$

hence by continuity

$$\lim_{\gamma \rightarrow 0} f^{-1}(\tau, \gamma) = \frac{\sqrt{\tau}}{(1-\sqrt{\tau})}.$$

Thus letting $\lambda = \lim_{\gamma \rightarrow 0} f^{-1}(\tau, \gamma)$, we get

$$\lim_{\gamma \rightarrow 0} \frac{\mathcal{E}_*(\tau, \gamma)}{\sigma^2 \gamma} = \lim_{\gamma \rightarrow 0} \int_{s \in [1-\sqrt{\gamma}, 1+\sqrt{\gamma}]} \frac{s}{(s+\lambda)^2} dH_\gamma(s) = \frac{1}{(1+\lambda)^2} = (1-\sqrt{\tau})^2.$$

By Eq. (28), we have that the limit equals the infimum

$$\lim_{\gamma \rightarrow 0} \frac{\mathcal{E}_*(\tau, \gamma)}{\sigma^2 \gamma} = \inf_{\gamma \in (0, \infty)} \frac{\mathcal{E}_*(\tau, \gamma)}{\sigma^2 \gamma}$$

which shows Eq. (29). □

From the above Theorem 2 we saw that $\mathcal{E}_*(\tau, \gamma)/\gamma \rightarrow (1-\sqrt{\tau})^2$ as $\gamma \rightarrow 0$, achieving its infimum in the limit. In the following theorem we will show that moreover $\mathcal{E}_*(\tau, \gamma)/\gamma$ is strictly increasing in γ which implies in particular that Eq. (28) becomes strictly looser as γ grows.

Theorem 3. *For fixed $\tau \in (0, 1]$, the ratio of the minimum excess error $\mathcal{E}_*(\tau, \gamma)$ to γ is increasing in γ , that is*

$$\frac{d\mathcal{E}_*(\tau, \gamma)/\gamma}{d\gamma} = \sigma^2 \frac{dE(f^{-1}(\tau, \gamma), \gamma)}{d\gamma} > 0, \quad \gamma \in (0, +\infty)$$

where the function E is defined in Eq. (25).

Proof. Let $\lambda = f^{-1}(\tau, \gamma)$. By the chain rule

$$\frac{dE(f^{-1}(\tau, \gamma), \gamma)}{d\gamma} = \frac{\partial E(\lambda, \gamma)}{\partial \gamma} + \frac{\partial E(\lambda, \gamma)}{\partial \lambda} \frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma}. \quad (30)$$

We will compute the following three terms in Eq. (30)

$$(1) : \frac{\partial E(\lambda, \gamma)}{\partial \gamma}, \quad (2) : \frac{\partial E(\lambda, \gamma)}{\partial \lambda}, \quad (3) : \frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma}.$$

The computation of the first two terms is direct

$$\begin{aligned} \frac{\partial E(\lambda, \gamma)}{\partial \gamma} &= \frac{\partial m(-\lambda; \gamma)}{\partial \gamma} - \lambda \frac{\partial m'(-\lambda; \gamma)}{\partial \gamma}, \\ \frac{\partial E(\lambda, \gamma)}{\partial \lambda} &= -2m'(-\lambda; \gamma) + \lambda m''(-\lambda; \gamma). \end{aligned}$$

For the third term, we take the derivative with respect to γ on both sides of Eq. (27) which yields

$$\frac{\partial f(\lambda, \gamma)}{\partial \lambda} \frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma} + \frac{\partial f(\lambda, \gamma)}{\partial \gamma} = 0.$$

Hence after re-arranging we have

$$\frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma} = -\frac{\frac{\partial f(\lambda, \gamma)}{\partial \gamma}}{\frac{\partial f(\lambda, \gamma)}{\partial \lambda}} = -\frac{\lambda^2 \frac{\partial m'(-\lambda; \gamma)}{\partial \gamma}}{2\lambda m'(-\lambda; \gamma) - \lambda^2 m''(-\lambda; \gamma)}.$$




Figure 2: For fixed $\gamma = n/p$ we plot the exact value of $\mathcal{E}_*(\tau)$ using Eq. (22) [solid] and the first-order Taylor approximation around $\tau = 0$ from Eq. 31 [dashed]. Observe that as $\gamma \rightarrow 1$ the approximation becomes less tight for larger values of τ .

Computing the product of terms (2) and (3) gives

$$\frac{\partial E(\lambda, \gamma)}{\partial \lambda} \frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma} = \lambda \frac{\partial m'(-\lambda; \gamma)}{\partial \gamma}.$$

Finally putting everything together in Eq. (30) yields

$$\begin{aligned} \frac{dE(f^{-1}(\tau, \gamma), \gamma)}{d\gamma} &= \frac{\partial E(\lambda, \gamma)}{\partial \gamma} + \frac{\partial E(\lambda, \gamma)}{\partial \lambda} \frac{\partial f^{-1}(\tau, \gamma)}{\partial \gamma} \\ &= \frac{\partial m(-\lambda; \gamma)}{\partial \gamma} - \lambda \frac{\partial m'(-\lambda; \gamma)}{\partial \gamma} + \lambda \frac{\partial m'(-\lambda; \gamma)}{\partial \gamma} \\ &= \frac{\partial m(-\lambda; \gamma)}{\partial \gamma} \Bigg|_{\lambda=f^{-1}(\tau, \gamma)} > 0, \end{aligned}$$

where the last inequality follows from Lemma 10. □

The previous bound provides a lower bound on the minimum excess risk that holds for all $\tau \in [0, 1]$ and becomes tight as $\gamma \rightarrow 0$ i.e. the amount of overparametrization becomes large. However, the bound is loose for small τ near the interpolation peak $\gamma = 1$. To understand behavior in this regime, for a fixed γ we compute a local expansion of $\mathcal{E}_*(\tau, \gamma)$ around $\mathcal{E}_*(0, \gamma)$. Instead of directly analyzing the function $\mathcal{E}_*(\tau, \gamma)$, we work with the function $\sqrt{\mathcal{E}_*(t, \gamma)}$ where we re-parametrize in terms of $t := \sqrt{\tau}$. We then obtain the bound by taking the first-order Taylor expansion of this function around $t = 0$ and then arguing it is a lower bound by showing the function is convex. In the end we convert this into lower bound on the original function $\mathcal{E}_*(\tau, \gamma)$. It is important to re-parametrize in terms of $\sqrt{\tau}$ since the derivative of $\mathcal{E}_*(\tau, \gamma)$ with respect to τ at 0 is $-\infty$, as is also true for the lower bound in Eq. (28). Intriguingly, as we explain in Remark 8, taking the Taylor expansion of the square-root of \mathcal{E}_* yields a tighter approximation rather than directly expanding \mathcal{E}_* .

Theorem 4 (MP lower bound for small τ). *For fixed $\gamma \in (0, 1)$ and $\tau \in [0, 1 - \gamma]$ we have*

$$\mathcal{E}_*(\tau, \gamma) \geq \frac{\sigma^2 \gamma}{1 - \gamma} \left(1 - \frac{\sqrt{\tau}}{\sqrt{1 - \gamma}} \right)^2. \quad (31)$$

Moreover, the lower bound is tight up to $o(\sqrt{\tau})$ error.

Proof. In the proof we will often suppress the dependence on γ and express dependence on τ in terms of $t := \sqrt{\tau}$. In particular, recalling the functions E , f , and f^{-1} defined in Eqs. (25), (26), and (27) we use

$$\begin{aligned}\mathcal{E}_*(t) &:= \mathcal{E}_*(\tau, \gamma), & f^{-1}(t) &:= f^{-1}(\tau, \gamma) \\ E(\lambda) &:= E(\lambda, \gamma), & f(\lambda) &:= f(\lambda, \gamma),\end{aligned}$$

and as a result Eq. (27) becomes

$$f(f^{-1}(t)) = t^2. \quad (32)$$

The first-order Taylor approximation of $\sqrt{\mathcal{E}_*(t)}$ around $t = 0$ is given by

$$P_1(t) := \sqrt{\mathcal{E}_*(0)} + t \cdot \left. \frac{d}{dt} \sqrt{\mathcal{E}_*(t)} \right|_{t=0} = \sqrt{\mathcal{E}_*(0)} + \frac{t}{2\sqrt{\mathcal{E}_*(0)}} \cdot \mathcal{E}'_*(0). \quad (33)$$

where as $t \rightarrow 0^+$

$$\sqrt{\mathcal{E}_*(t)} = P_1(t) + o(t). \quad (34)$$

Note that since $\lim_{t \rightarrow 0^+} f^{-1}(t, \gamma) = 0$ for any γ , by Lemma 8 we

$$\mathcal{E}'_*(0) = \sigma^2 \gamma \lim_{\lambda \rightarrow 0^+} [m(-\lambda; \gamma) - \lambda m'(-\lambda; \gamma)] = \sigma^2 \gamma \lim_{\lambda \rightarrow 0^+} m(-\lambda; \gamma) = \frac{\sigma^2 \gamma}{1 - \gamma}.$$

Furthermore letting $\lambda = f^{-1}(t)$, we have by the chain rule

$$\mathcal{E}'_*(t) = \sigma^2 \gamma \frac{dE(f^{-1}(t))}{dt} = \sigma^2 \gamma \frac{dE(\lambda)}{d\lambda} \frac{df^{-1}(t)}{dt}. \quad (35)$$

We can directly calculate the first term

$$\frac{dE(\lambda)}{d\lambda} = -2m(-\lambda; \gamma) + \lambda m''(-\lambda; \gamma).$$

By differentiating both sides of Eq. (32) with respect to t we get

$$\frac{df^{-1}(t)}{dt} \frac{df(\lambda)}{d\lambda} = 2t.$$

Therefore rearranging the above gives

$$\begin{aligned}\frac{df^{-1}(t)}{dt} &= \frac{2t}{f'(\lambda)} = \frac{2t}{2\lambda m'(-\lambda; \gamma) - \lambda^2 m''(-\lambda; \gamma)} \\ &= \frac{2\lambda \sqrt{m'(-\lambda; \gamma)}}{2\lambda m'(-\lambda; \gamma) - \lambda^2 m''(-\lambda; \gamma)} \\ &= \frac{2\sqrt{m'(-\lambda; \gamma)}}{2m'(-\lambda; \gamma) - \lambda m''(-\lambda; \gamma)}.\end{aligned}$$

Hence from Eq. (35) we see

$$\mathcal{E}'_*(t) = -2\sigma^2 \gamma \sqrt{m'(-\lambda; \gamma)}. \quad (36)$$

Observe that since $m'(-\lambda; \gamma) \geq 0$, the derivative $\mathcal{E}'_*(t) \leq 0$. Furthermore, since λ is increasing in t and $m'(-\lambda; \gamma)$ is decreasing in λ , it follows that $m'(-\lambda; \gamma)$ is decreasing in t . Thus $\mathcal{E}'_*(t)$ is increasing, i.e. $\mathcal{E}''_*(t) \geq 0$. Using Lemma 9

$$\mathcal{E}'_*(0) = \lim_{\lambda \rightarrow 0^+} -2\sigma^2 \gamma \sqrt{m'(-\lambda; \gamma)} = -2\gamma \sigma^2 \frac{1}{(1 - \gamma)^{3/2}}.$$

Plugging this into the Taylor approximation $P_1(t)$ in Eq. (33) we get that

$$\begin{aligned} P_1(t) &= \sigma \sqrt{\frac{\gamma}{1-\gamma}} - \sigma t \sqrt{\frac{1-\gamma}{\gamma}} \frac{\gamma}{(1-\gamma)^{3/2}} \\ &= \sigma \sqrt{\frac{\gamma}{1-\gamma}} \left(1 - t \frac{1}{\sqrt{1-\gamma}}\right). \end{aligned}$$

We now show that $\sqrt{\mathcal{E}_*(t)} \geq P_1(t)$ by showing that $\sqrt{\mathcal{E}_*(t)}$ is convex on $[0, 1]$ for which it suffices to show that the second derivative is non-negative. Computing the second derivative, we get

$$\frac{d^2}{dt^2} \sqrt{\mathcal{E}_*(t)} = \frac{\mathcal{E}_*(t)}{2\sqrt{\mathcal{E}_*(t)}} - \frac{\mathcal{E}'_*(t)}{4\mathcal{E}_*(t)^{3/2}}.$$

We can see that this is always non-negative since $\mathcal{E}_*(t) \geq 0$ and from Eq. (36) and the accompanying remarks $\mathcal{E}'_*(t) \leq 0$ and $\mathcal{E}''_*(t) \geq 0$. Returning to the original τ parameterization, we have that for all $\tau \in [0, 1]$

$$\sqrt{\mathcal{E}_*(\tau)} \geq P_1(\sqrt{\tau}) = \sigma \sqrt{\frac{\gamma}{1-\gamma}} \left(1 - \sqrt{\frac{\tau}{1-\gamma}}\right).$$

Note that for $\tau \leq 1 - \gamma$, both sides of the above inequality are non-negative. Hence we can square both sides yielding

$$\mathcal{E}_*(\tau, \gamma) \geq \sigma^2 \frac{\gamma}{1-\gamma} \left(1 - \sqrt{\frac{\tau}{1-\gamma}}\right)^2$$

which was the desired lower bound. The fact that the bound is tight up to $o(\sqrt{\tau})$ follows from Eq. (34). \square

Remark 8. Observe that from the above computations the first-order expansion of $\mathcal{E}_*(t)$ is given by

$$\mathcal{E}_*(0) + \sqrt{\tau} \mathcal{E}'_*(0) = \sigma^2 \sqrt{\frac{\gamma}{1-\gamma}} \left(1 - 2\sqrt{\frac{\tau}{1-\gamma}}\right) < \sigma^2 \frac{\gamma}{1-\gamma} \left(1 - \sqrt{\frac{\tau}{1-\gamma}}\right)^2 \leq \mathcal{E}_*(\tau)$$

which shows that expanding the square-root of the minimum excess error yields a tighter bound.

Note that the bound in Eq. (31) matches the lower bound in Theorem 2 as $\gamma \rightarrow 0$. Interestingly, this bound has the form of a multiplicative factor of the minimum interpolation ($\tau = 0$) loss and is only valid for small τ . We will in fact show that no bound of that form can capture the behavior of the excess loss for arbitrarily small $\tau > 0$ as $\gamma \rightarrow 1$ in Theorem 5. In particular, we now show a discontinuity at $\tau = 0$ when considering the peak $n = p$.

Theorem 5 (Excess Loss at Peak). For $\gamma = 1$ and $\tau \in (0, 1)$,

$$\frac{\mathcal{E}_*(\tau, 1)}{\sigma^2} = \frac{1}{4\tau} + \frac{\tau}{4} - \frac{1}{2}$$

hence as $\tau \rightarrow 0$,

$$\mathcal{E}_*(\tau, 1) = \Theta(\tau^{-1}).$$

Proof. Consider the Stieltjes transform $m(-z; \gamma)$ and its derivative. At the interpolation peak $\gamma = 1$, we have

$$\begin{aligned} m(-z; 1) &= \frac{1}{2} \left(\sqrt{1 + 4/z} - 1 \right) \\ m'(-z; 1) &= -\frac{z+2}{2z\sqrt{z^2+4z}} + \frac{\sqrt{z^2+4z}}{2z^2}. \end{aligned}$$

One can then check that

$$z^2 m'(-z; 1) = \frac{1}{\sqrt{1 + 4/z}}.$$

Hence the fixed point equation Eq. (23) can be written as

$$\tau = \lambda^2 m'(-\lambda; 1) = \frac{1}{\sqrt{1 + 4/\lambda}}$$

which implies that λ satisfies the following

$$\begin{aligned} \sqrt{1 + 4/\lambda} &= \frac{1}{\tau}, \\ \frac{1}{\lambda} &= \frac{1}{4} \left(\frac{1}{\tau^2} - 1 \right). \end{aligned}$$

Using the above relations, we can compute the minimum excess error using Eq. (22)

$$\begin{aligned} \frac{\mathcal{E}_*(\tau, 1)}{\sigma^2} &= m(-\lambda; 1) - \lambda m'(-\lambda; 1) \\ &= \frac{1}{2} \left(\sqrt{1 + 4/\lambda} - 1 \right) - \frac{\tau}{\lambda} \\ &= \frac{1}{2} \left(\frac{1}{\tau} - 1 \right) - \frac{1}{4} \left(\frac{1}{\tau} - \tau \right) \\ &= \frac{1}{4\tau} + \frac{\tau}{4} - \frac{1}{2}. \end{aligned}$$

□

Remark 9. As mentioned earlier, a consequence of the above is that we cannot have a lower bound of the form $\mathcal{E}_*(\tau; n, p) \geq \mathcal{E}_*(0; n, p) \cdot f(\tau)$ for some function f which satisfies $f(\tau) > 0$ for $\tau > 0$ since for any $\tau > 0$,

$$\lim_{\gamma \rightarrow 1} \frac{\mathcal{E}_*(\tau, \gamma)}{\mathcal{E}_*(0, \gamma)} = 0.$$

4 Conclusion

In this work we demonstrated a trade-off between the expected loss, empirical loss, and the number of parameters for general linear models. In particular we have shown that near-optimal algorithms output models that are either classical (with empirical loss approaching the noise level) or have significant excess over-parameterization, i.e., have many more parameters than the number needed to fit the training data. This trade-off is universal as it is non-asymptotic, holds for any algorithm, and any data distribution (under mild non-degeneracy assumptions).

We also provided a more precise asymptotic lower bound under Marchenko-Pastur distributional assumptions near the classical double descent peak where the amount of overparametrization is just enough to interpolate the data. Remarkably however, as the level of overparametrization increases the minimum excess loss exactly matches the universal bound, demonstrating the tightness of the bound.

The open questions that remain include extending our results to more general non-linear parametric families and to classification settings.

Acknowledgements

The authors would like to thank Amirhesam Abedsoltan for finding an error in a previous version of the proof of Theorem 1. Correcting the proof led to an improved lower bound which is now tight. We also thank the anonymous reviewers for insightful comments. We are grateful for support from the National Science Foundation (NSF) and the Simons Foundation for the Collaboration on the Theoretical Foundations of Deep Learning³ through awards DMS-2031883 and #814639 as well as NSF IIS-1815697 and the TILOS institute (NSF CCF-2112665). NG would also like to acknowledge support from the NSF RTG Grant #1745640.

³<https://deepfoundations.ai/>

References

- [1] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- [2] Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1):67–83, 2020.
- [3] Luc Devroye, Laszlo Györfi, and Adam Krzyżak. The hilbert kernel regression estimate. *Journal of Multivariate Analysis*, 65(2):209–227, 1998.
- [4] Mikhail Belkin, Daniel J Hsu, and Partha Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. *Advances in neural information processing systems*, 31, 2018.
- [5] Mikhail Belkin, Alexander Rakhlin, and Alexandre B Tsybakov. Does data interpolation contradict statistical optimality? In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1611–1619. PMLR, 2019.
- [6] Julien Chhor, Suzanne Sigalla, and Alexandre B Tsybakov. Benign overfitting and adaptive nonparametric regression. *arXiv preprint arXiv:2206.13347*, 2022.
- [7] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986, 2022.
- [8] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- [9] Frederic Koehler, Lijia Zhou, Danica J Sutherland, and Nathan Srebro. Uniform convergence of interpolators: Gaussian width, norm bounds and benign overfitting. *Advances in Neural Information Processing Systems*, 34:20657–20668, 2021.
- [10] Alexander Tsigler and Peter L Bartlett. Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*, 2020.
- [11] Niladri S Chatterji and Philip M Long. Foolish crowds support benign overfitting. *Journal of Machine Learning Research*, 23(125):1–12, 2022.
- [12] Difan Zou, Jingfeng Wu, Vladimir Braverman, Quanquan Gu, and Sham Kakade. Benign overfitting of constant-size sgd for linear regression. In *Conference on Learning Theory*, pages 4633–4635. PMLR, 2021.
- [13] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.
- [14] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Generalization error of random feature and kernel methods: hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59:3–84, 2022.
- [15] Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *The annals of statistics*, 2020.
- [16] Ben Adlam and Jeffrey Pennington. The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization. In *International Conference on Machine Learning*, pages 74–84. PMLR, 2020.
- [17] Yuan Cao, Zixiang Chen, Mikhail Belkin, and Quanquan Gu. Benign overfitting in two-layer convolutional neural networks. *arXiv preprint arXiv:2202.06526*, 2022.
- [18] Spencer Frei, Niladri S Chatterji, and Peter L Bartlett. Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. *arXiv preprint arXiv:2202.05928*, 2022.
- [19] Spencer Frei, Gal Vardi, Peter L Bartlett, and Nathan Srebro. Benign overfitting in linear classifiers and leaky relu networks from kkt conditions for margin maximization. *arXiv preprint arXiv:2303.01462*, 2023.

- [20] Zhu Li, Zhi-Hua Zhou, and Arthur Gretton. Towards an understanding of benign overfitting in neural networks. *arXiv preprint arXiv:2106.03212*, 2021.
- [21] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [22] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [23] David Holzmüller. On the universality of the double descent peak in ridgeless regression. *arXiv preprint arXiv:2010.01851*, 2020.
- [24] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- [25] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009.
- [26] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- [27] Chen Cheng, John Duchi, and Rohith Kuditipudi. Memorize to generalize: on the necessity of interpolation in high dimensional linear regression. *arXiv preprint arXiv:2202.09889*, 2022.
- [28] Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–959, 2020.
- [29] Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 123–132, 2021.
- [30] Sebastien Bubeck and Mark Sellke. A universal law of robustness via isoperimetry. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 28811–28822. Curran Associates, Inc., 2021.
- [31] Lee H Dicker. Ridge regression and asymptotic minimax estimation over spheres of growing dimension. 2016.
- [32] Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.
- [33] Godfrey Harold Hardy, John Edensor Littlewood, George Pólya, György Pólya, et al. *Inequalities*. Cambridge university press, 1952.

A Missing Proofs from Section 2.1

In this section of the appendix, we supply missing proofs for some the auxiliary results used in the proof of Theorem 1 in Section 2.1. Recall that we denote the optimal linear predictor as $\beta_\star := \arg \min_{\beta \in \mathbb{R}^p} R(\beta)$.

Lemma 5 (Optimal Linear Predictor). *Define β_\star as in Eq. (5). Then*

1. β_\star is the orthogonal projection of f_\star onto the subspace of linear functions $\mathcal{H} = \{\beta(\mathbf{x}) : \beta \in \mathbb{R}^p\}$ in L^2 ,
2. $\mathcal{E}(f) := R(f) - R(f_\star) = \mathbb{E}_{\mathbf{x}}[(f(\mathbf{x}) - f_\star(\mathbf{x}))^2]$,
3. $\mathcal{E}^{\text{lin}}(\beta) := R(\beta) - R(\beta_\star) = \mathbb{E}_{\mathbf{x}}[(\beta(\mathbf{x}) - \beta_\star(\mathbf{x}))^2]$.

Proof. By definition of the optimal linear predictor

$$\begin{aligned}
 \beta_\star &= \arg \min_{\beta \in \mathcal{H}} R(\beta) \\
 &= \arg \min_{\beta \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y)}[(y - \beta(\mathbf{x}))^2] \\
 &= \arg \min_{\beta \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y)}[(y - f_\star(\mathbf{x}))^2 + (f_\star(\mathbf{x}) - \beta(\mathbf{x}))^2] \quad (\text{since } \mathbb{E}_{y|\mathbf{x}}[y - f_\star(\mathbf{x})] = 0) \\
 &= \arg \min_{\beta \in \mathcal{H}} \mathbb{E}_{\mathbf{x}}[(f_\star(\mathbf{x}) - \beta(\mathbf{x}))^2],
 \end{aligned}$$

which shows Claim 1. For Claim 2

$$\begin{aligned}
 R(f) - R(f_\star) &= \mathbb{E}_{(\mathbf{x}, y)}[(y - f(\mathbf{x}))^2] - \mathbb{E}_{(\mathbf{x}, y)}[(y - f_\star(\mathbf{x}))^2] \\
 &= \mathbb{E}_{(\mathbf{x}, y)}[(f(\mathbf{x}) - f_\star(\mathbf{x}))^2] - 2\mathbb{E}_{(\mathbf{x}, y)}[(f(\mathbf{x}) - f_\star(\mathbf{x}))(y - f_\star(\mathbf{x}))] \\
 &= \mathbb{E}_{\mathbf{x}}[(f(\mathbf{x}) - f_\star(\mathbf{x}))^2].
 \end{aligned}$$

For Claim 3

$$\begin{aligned}
 R(\beta) - R(\beta_\star) &= \mathcal{E}(\beta) - \mathcal{E}(\beta_\star) \\
 &= \mathbb{E}_{\mathbf{x}}[(f_\star(\mathbf{x}) - \beta(\mathbf{x}))^2 - (f_\star(\mathbf{x}) - \beta_\star(\mathbf{x}))^2] \\
 &= \mathbb{E}_{\mathbf{x}}[(\beta(\mathbf{x}) - \beta_\star(\mathbf{x}))^2] - 2\mathbb{E}_{\mathbf{x}}[(f_\star(\mathbf{x}) - \beta_\star(\mathbf{x}))(\beta(\mathbf{x}) - \beta_\star(\mathbf{x}))] \\
 &= \mathbb{E}_{\mathbf{x}}[(\beta(\mathbf{x}) - \beta_\star(\mathbf{x}))^2]
 \end{aligned}$$

where the second equality holds from Claim 2 and the last equality holds from Claim 1 in Lemma 5. □

Lemma 6 (Excess Linear Loss). *The excess loss satisfies the following lower bound*

$$\mathcal{E}(\beta) \geq \mathcal{E}^{\text{lin}}(\beta) = \left\| \Sigma^{1/2}(\beta - \beta_\star) \right\|^2.$$

Proof. It is clear that $\mathcal{E}(\beta) \geq \mathcal{E}^{\text{lin}}(\beta)$. Using Lemma 5 and the definition of Σ

$$\begin{aligned}
 \mathcal{E}(\beta) &= \mathbb{E}_{\mathbf{x}}(\beta(\mathbf{x}) - \beta_\star(\mathbf{x}))^2 \\
 &= \mathbb{E}_{\mathbf{x}}(\beta^\top \phi(\mathbf{x}) - \beta_\star^\top \phi(\mathbf{x}))^2 \\
 &= (\beta - \beta_\star)^\top \mathbb{E}_{\mathbf{x}}[\phi(\mathbf{x})\phi(\mathbf{x})^\top](\beta - \beta_\star) \\
 &= \left\| \Sigma^{1/2}(\beta - \beta_\star) \right\|^2.
 \end{aligned}$$

□

Recall that we defined the random vectors

$$\boldsymbol{\xi}_{\text{noise}} = (y_i - f_\star(\mathbf{x}_i))_{i \in [n]} \in \mathbb{R}^n, \quad \boldsymbol{\xi}_{\text{lin}} = (f_\star(\mathbf{x}_i) - \beta_\star(\mathbf{x}_i))_{i \in [n]} \in \mathbb{R}^n$$

and let $\boldsymbol{\xi} = \boldsymbol{\xi}_{\text{noise}} + \boldsymbol{\xi}_{\text{lin}}$.

Lemma 7 (Expectation Over Noise). *Let $f : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}_+^n$ be any PSD matrix valued function. Then,*

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}}[\boldsymbol{\xi}^\top f(\mathbf{X}) \boldsymbol{\xi}] \geq \sigma^2 \mathbb{E}_{\mathbf{X}} \text{Tr}(f(\mathbf{X})). \quad (37)$$

Moreover, if $\text{Var}(y | \mathbf{x}) = \sigma^2$ almost surely and $f_\star = \beta_\star$ then the above is an equality.

Proof. Recalling that $\boldsymbol{\xi} = \boldsymbol{\xi}_{\text{noise}} + \boldsymbol{\xi}_{\text{lin}}$. By definition of $\boldsymbol{\xi}_{\text{noise}}$ we have that

$$\mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}}] = \mathbf{0}$$

and by the assumption on the noise variance we have almost surely

$$\mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}} \boldsymbol{\xi}_{\text{noise}}^\top] = \text{diag}((\text{Var}[y | \mathbf{x}_i])_{i \in [n]}) \succeq \sigma^2 \mathbf{I}. \quad (38)$$

Using these observations we have

$$\begin{aligned} \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}^\top f(\mathbf{X}) \boldsymbol{\xi}] &= \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{noise}}] + 2\mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{noise}}] + \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{lin}}] \\ &= \text{Tr}\left(f(\mathbf{X}) \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}} \boldsymbol{\xi}_{\text{noise}}^\top]\right) + 2\boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}}] + \boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{lin}} \\ &= \text{Tr}\left(f(\mathbf{X}) \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}_{\text{noise}} \boldsymbol{\xi}_{\text{noise}}^\top]\right) + \boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{lin}} \\ &\geq \sigma^2 \text{Tr}(f(\mathbf{X})) + \boldsymbol{\xi}_{\text{lin}}^\top f(\mathbf{X}) \boldsymbol{\xi}_{\text{lin}} \\ &\geq \sigma^2 \text{Tr}(f(\mathbf{X})), \end{aligned}$$

where we used Eq. (38) in the first inequality and we used the assumption that $f(\mathbf{X})$ is a PSD matrix for the last inequality. By iterating expectations we have

$$\mathbb{E}_{\mathbf{y}, \mathbf{X}}[\boldsymbol{\xi}^\top f(\mathbf{X}) \boldsymbol{\xi}] = \mathbb{E}_{\mathbf{X}} \mathbb{E}_{\mathbf{y} | \mathbf{X}}[\boldsymbol{\xi}^\top f(\mathbf{X}) \boldsymbol{\xi}] \geq \sigma^2 \mathbb{E}_{\mathbf{X}} \text{Tr}(f(\mathbf{X}))$$

as desired. Note that if $\text{Var}(y | \mathbf{x}) = \sigma^2$ then Eq. (38) becomes an equality and if $f_\star = \beta_\star$ then $\boldsymbol{\xi}_{\text{lin}} = \mathbf{0}$ and so it easy to see that Eq. (37) is an equality as well. \square

B Marchenko-Pastur Law

In this section, we will give some definitions, facts, and results concerning the Marchenko-Pastur Law.

Definition 1 (Marchenko-Pastur Law). *The Marchenko-Pastur Law with aspect ratio γ , denoted $\text{MP}(\gamma)$, is given the probability distribution*

$$\left(1 - \frac{1}{\gamma}\right)_+ \delta_0 + \frac{1}{\gamma 2\pi x} \sqrt{(b-x)(x-a)}$$

where

$$a = (1 - \sqrt{\gamma})^2, \quad b = (1 + \sqrt{\gamma})^2.$$

Definition 2 (Stieltjes Transform). *For $z > 0$ and $\gamma > 0$, the Stieltjes transform of the Marchenko-Pastur law is*

$$m(-z; \gamma) := \int \frac{1}{s+z} d\text{MP}_\gamma(s) = \frac{-(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma z}$$

and its derivative satisfies

$$\begin{aligned} m'(-z; \gamma) &= \int \frac{1}{(s+z)^2} d\text{MP}_\gamma(s) \\ &= \frac{-(1-\gamma+z) - 2\gamma}{2\gamma z \sqrt{(1-\gamma+z)^2 + 4\gamma z}} + \frac{1}{2\gamma z} + \frac{\sqrt{(1-\gamma+z)^2 + 4\gamma z} - (1-\gamma+z)}{2z^2\gamma}. \end{aligned}$$

Lemma 8. *For any $\gamma > 0$,*

$$\lim_{z \rightarrow 0^+} m(-z; \gamma) = \frac{1}{1-\gamma}$$

Proof. This just follows from a direct computation

$$\begin{aligned} \lim_{z \rightarrow 0^+} m(-z; \gamma) &= \lim_{z \rightarrow 0^+} \frac{-(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma z} \\ &= \lim_{z \rightarrow 0^+} \frac{-(1-\gamma+z)^2 + (1-\gamma+z)^2 + 4\gamma z}{2\gamma z[(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}]} \\ &= \lim_{z \rightarrow 0^+} \frac{2}{(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}} = \frac{1}{1-\gamma}. \end{aligned}$$

□

Lemma 9. *For any $\gamma > 0$,*

$$\lim_{z \rightarrow 0^+} m'(-z; \gamma) = \frac{1}{(1-\gamma)^3}$$

Proof. By the definition of derivative

$$\begin{aligned} \lim_{z \rightarrow 0^+} m'(-z; \gamma) &= \lim_{z \rightarrow 0^+} \frac{m(-z; \gamma) - m(0; \gamma)}{-z} \\ &= \lim_{z \rightarrow 0^+} \frac{1}{z} \left(\frac{(1-\gamma+z) - \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma z} + \frac{1}{1-\gamma} \right) \\ &= \lim_{z \rightarrow 0^+} \frac{(1-\gamma)(1-\gamma+z) - (1-\gamma)\sqrt{(1-\gamma+z)^2 + 4\gamma z} + 2\gamma z}{2\gamma(1-\gamma)z^2} \\ &= \lim_{z \rightarrow 0^+} \frac{(1-\gamma)^2 + z(1+\gamma) - (1-\gamma)\sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma(1-\gamma)z^2} \\ &= \lim_{z \rightarrow 0^+} \frac{[(1-\gamma)^2 + z(1+\gamma)]^2 - (1-\gamma)^2[(1-\gamma+z)^2 + 4\gamma z]}{2\gamma(1-\gamma)z^2[(1-\gamma)^2 + z(1+\gamma) + (1-\gamma)\sqrt{(1-\gamma+z)^2 + 4\gamma z}]} \end{aligned}$$

Observe that the numerator simplifies as

$$\begin{aligned}
& [(1-\gamma)^2 + z(1+\gamma)]^2 - (1-\gamma)^2[(1-\gamma+z)^2 + 4\gamma z] \\
&= (1-\gamma)^4 + 2z(1-\gamma)^2(1+\gamma) + z^2(1+\gamma)^2 - [(1-\gamma)^4 + 2z(1-\gamma)^2(1+\gamma) + (1-\gamma)^2 z^2] \\
&= z^2[(1+\gamma)^2 - (1-\gamma)^2] = 4z^2\gamma.
\end{aligned}$$

Hence

$$\begin{aligned}
\lim_{z \rightarrow 0^+} m'(-z; \gamma) &= \lim_{z \rightarrow 0^+} \frac{4\gamma}{2\gamma(1-\gamma)[(1-\gamma)^2 + z(1+\gamma) + (1-\gamma)\sqrt{(1-\gamma+z)^2 + 4\gamma z}]} \\
&= \frac{2}{(1-\gamma)[(1-\gamma)^2 + (1-\gamma)\sqrt{(1-\gamma)^2}]} = \frac{1}{(1-\gamma)^3}.
\end{aligned}$$

□

Lemma 10. For any $z > 0$,

$$\frac{\partial m(-z; \gamma)}{\partial \gamma} > 0.$$

Proof. We can rewrite the Stieltjes transform as follows

$$\begin{aligned}
m(-z; \gamma) &= \frac{-(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma z} \\
&= \frac{-(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{2\gamma z} \frac{(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}{(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}} \\
&= \frac{2}{(1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}}.
\end{aligned}$$

Therefore it suffices to show that

$$f(\gamma) = (1-\gamma+z) + \sqrt{(1-\gamma+z)^2 + 4\gamma z}$$

is decreasing in γ . Taking the derivative we see that

$$f'(\gamma) = -1 + \frac{-(1-\gamma+z) + 2z}{\sqrt{(1-\gamma+z)^2 + 4\gamma z}}.$$

Note that $f'(\gamma) < 0$ if and only if

$$2z - (1-\gamma+z) < \sqrt{(1-\gamma+z)^2 + 4\gamma z}.$$

Squaring both sides and clearing terms this is equivalent to

$$4z^2 - 4z(1-\gamma+z) < 4\gamma z.$$

Using the fact that $z > 0$ we can divide both sides by $4z$ to get

$$z - (1-\gamma+z) < \gamma$$

which is clearly true. □

C Auxiliary Lemmas

Lemma 11 (AM-HM inequality). *Let $x_1, \dots, x_n \in \mathbb{R}$. Then*

$$\sum_{i=1}^n \frac{1}{x_i} \geq n^2 \left(\sum_{i=1}^n x_i \right)^{-1}.$$

Proof. By the Cauchy-Schwarz inequality

$$n^2 = \left(\sum_{i=1}^n \sqrt{x_i} \cdot \frac{1}{\sqrt{x_i}} \right)^2 \leq \sum_{i=1}^n \sqrt{x_i}^2 \cdot \sum_{i=1}^n \frac{1}{\sqrt{x_i}^2},$$

which after re-arranging gives the desired inequality. □

Lemma 12 (Chebyshev's Sum Inequality [33]). *If $a_1 \leq a_2 \leq \dots \leq a_n$ and $b_1 \geq \dots \geq b_n$ then*

$$\frac{1}{n} \sum_{i=1}^n a_i b_i \leq \left(\frac{1}{n} \sum_{i=1}^n a_i \right) \left(\frac{1}{n} \sum_{i=1}^n b_i \right).$$

Lemma 13 (Concentration of Quadratic Forms, Lemma 7.6 from [32]). *For each positive integer n , let \mathbf{A}_n be a random $n \times n$ PSD matrix. Let ξ_1, ξ_2, \dots be a sequence of i.i.d random variables and denote $\boldsymbol{\xi}_n = (\xi_1, \dots, \xi_n)$. Then if*

- $\sup_n \|\mathbf{A}_n\| = O(1)$
- $\mathbb{E}[\xi_1] = 0$, $\mathbb{E}[\xi_1^2] = \sigma^2$, and $\mathbb{E}[\xi_1^{4+\eta}]$ for some $\eta > 0$

we have the following convergence

$$\frac{1}{n} \boldsymbol{\xi}_n^\top \mathbf{A}_n \boldsymbol{\xi}_n - \frac{\sigma^2}{n} \text{Tr}(\mathbf{A}_n) \rightarrow 0$$

almost surely.